



# WLCG Collaboration Workshop

Welcome & Introduction

CHEP 2007

Victoria, BC, Canada

# Introduction

- 2<sup>nd</sup> WLCG Collaboration workshop
    - 1<sup>st</sup> was held at CERN January 2006
  - Follows several “Service Challenge” workshops
    - CERN, Mumbai, ...
  - Each with a specific focus...
- **Focus of this workshop:** Dress Rehearsals, Cosmics Runs and Readiness of Residual Services
- **Robust & resilient services is another important issue – see Operations / Monitoring BOFs...**
  - **Plan a follow-on workshop on this topic in November**
    - IT amphitheatre booked week of November 26<sup>th</sup>
  - **Next WLCG Collaboration workshop – April 21 – 25 2008, CERN**
  - **Also pre-CHEP 2009 – March 21 – 22 2009, Prague**
  - **Grid Camp 2007, WLCG Tier2 & EU-IndiaGrid Workshop, Oct 28 – Nov 2, Taipei**

# ISSeG Training Session: Improving Site Security

- Tuesday 04 September 2007
- 14:00 – 18:00 (Room: Oak Bay)
- The topics to be covered are:
  - How to sell security – and get resources for it
  - Risk assessments
  - Recommendations
  - Feedback/Questions
- There will be no charge for attending the session.



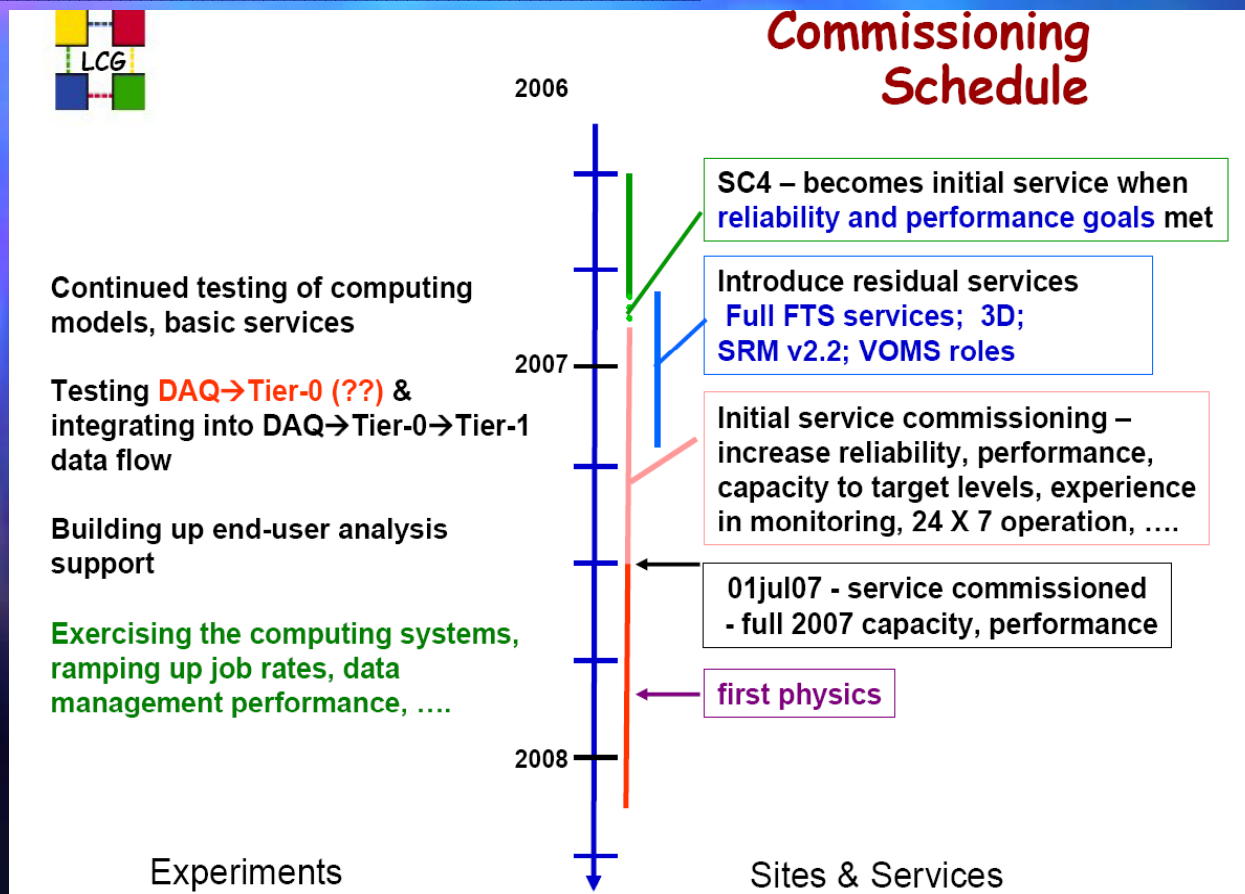
## Workshop Goals

- To understand the state of *{site, experiment, service}* readiness for data taking *{cosmics, pp collisions}* according to the current LHC schedule;
- To identify the key outstanding issues and associated milestones;
- To understand experiment activities *{'dress rehearsals', cosmics runs in November/December 2007, March 2008, pp from Spring/Summer on}*.

# Workshop Agenda

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

# 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



## Quality of Services - July GDB

- More than one 'thread' in this area:
  - Monitoring working groups (monitoring BOF)
    - **"Reliability by design"**
      - see [Stockholm workshop](#), operations BOF and [CHEP talk](#)
    - CMS Critical Services review.
- **"Resilience to glitches" is possible & (highly) desirable**
- (Re-)reviewing CERN services to understand current status & steps required to reach goals
- **I believe that this is also possible (at cost) for storage services - to be addressed as a second step (in the post-SRM v2.2 world?)**
- *"pro-active: (of a policy or person or action) controlling a situation by causing something to happen rather than waiting to respond to it after it happens"*

# Building Robust Services

- In 2005, did an *“a priori”* analysis of main services and their criticality
- This led to a deployment plan, including h/w requirements, operations issues etc.
- We have now quite extensive experience running production services, and hence an *“a posteriori”* analysis is called for
- **But this must be a coordinated, end-to-end analysis, and not just individual (component) services**
- An analogy (oh no...) follows...





# WLCG Tier1 Services<sup>1</sup>

- 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.

<sup>1</sup> WLCG Memorandum of Understanding (signed by each T0/T1/T2)

## 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
	<a href="#">PhEDEx</a>	Stops all transfers between sites for all CMS	Critical Service
CMS RB and BDII	Used by <a href="#">CRAB</a> and <a href="#">ProdAgent</a> 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 CERN to Tier-1s fail. There is a multi-day output buffer at the Tier-0 and the networking requirements have a factor of 2 headroom for recovery	
	Used to send data to and from Russian Tier-2 sites	Simulation processing at Russian Tier-2s will remain in the export buffers. The computing model calls for multi-day buffer and data could be temporarily archived at an alternate Tier-1	
SRM at CERN	Used by CERN transfers to and from Tier-1s	Transfers from CERN to Tier-1s fail. There is a multi-day output buffer at the Tier-0 and the networking requirements have a factor of 2 headroom for recovery	
	Used to send data to and from Russian Tier-2 sites	Simulation processing at Russian Tier-2s will remain in the export buffers. The computing model calls for multi-day buffer and data could be temporarily archived at an alternate Tier-1	

Castor at CERN	Prompt Reconstruction, Custodial Data Archive Used by calibration users, CMS Analysis Facility work.	The input buffer in the Tier-0 has a 1-2 day capacity. In ability to get data out of Castor impacts the ability to perform calibration and alignment which impacts the quality of prompt reconstruction	Critical Service
LFC	Currently used by <a href="#">DLS</a>	Prevents new analysis and production jobs from resolving the location of data. Being replaced with a dependency on the Oracle DB in the next version of the CMS data management components	
SAM	Currently used for site monitors	Used for information purposes currently, but will shortly be used as a requirement for site submission of organized processing jobs	Critical Service if used in job submission chain
GOCDDB	Used to discover downtimes for SAM tests	Used for information purposes currently, but will shortly be used as a requirement for site submission of organized processing	
dashboard	Used for system monitoring. Relies on Oracle	Used for information purposes and local site monitoring and alarm systems	
CMS Services			
Dataset Bookkeeping Server	Used by analysis and production job specification	The new Tier-0 prompt reconstruction workflow relies on DBS and analysis jobs through <a href="#">CRAB</a> rely on DBS for resolving the logical files names of a dataset. Proposal is to introduce local scope DBS systems that would rely on local MySQL DBs to ensure production running will continue with break in central services	
<a href="#">PhEDEx</a> instance at CERN	Used for staging simulation to CERN	Impacts the ability to bring new simulation samples for limited CAF analysis and the ability to archive simulation from the Russian Tier-2 centers.	
<a href="#">PhEDEx</a> Monitoring Pages	Used by some sites for monitoring and alarm generation	Transfers continue to succeed, but its difficult to see the performance	
Frontier Launchpad	Used by Frontier system to serve database caches	Impacts the ability to update caches.	



Operations BOF this afternoon...



## Critical Service Review - Status of the T0 services



## Operations BOF

- **Presentation on the review of Tier0 services wrt robustness / resilience**
- Some further suggestions raised at the WLCG Management Board are listed below:
- *H.Marten suggested that ... the sites should have the top 5 issues for each service.*
- *I.Fisk also suggested a session for feedback from the sites to the developers, in order to improve the services and the software developed.*
- *R.Tafirout suggested that sites should share solutions; sometimes the information is known by others but not easily accessible.*
- *I.Fisk asked for the correlation between “availability as seen via SAM” vs. the “real jobs success rates of the Experiments (via Crab, Ganga, etc)”.*



## DM BOF(s)

1. Update on 'the plan' for SRM v2.2 production roll-out: have we met interim milestones for the summer? Up-coming milestones for production deployment in two phases:
  - i. Early adopters post-CSA07 (Oct 15<sup>th</sup> on);
    - Assuming that implementations are ready for production release, including adequate documentation and experiment testing
  - ii. Remaining WLCG sites < end 2007. (US Tier2s Feb 2008)
2. Tutorial / training on site setup and configuration according to LHC experiments' requirements. Experience from early adopters / beta-testers from SRM v2.2 and related deployment.
3. Other issues (e.g. compatibility of SRM implementations)
4. Planning processes - difference between Europe & US, e.g. for dCache deployment





## DB BOF

- Tier 0 / Tier 1 database requirement update
  - summary for ATLAS, CMS and LHCb
- Results from the ATLAS and LHCb scalability tests against Tier 1 database replicas
- Status of LFC back-end database installations at LHCb Tier 1 sites
- CMS Frontier Deployment - Status update



## User Support

- The aim is to better understand what the VOs expect from the WLCG/EGEE support infrastructure and how they interact with it, in order to improve the processes and workflows and to make the EGEE user support more helpful for the VO-users.
  - This is a follow-up and continuation of the User Support Workshop held in Karlsruhe in March 2007.
- The latest developments in the EGEE/WLCG user support will be presented.
- Updates from the VOs about their support infrastructure are of course welcome.



## Monitoring BOF

- Progress since last workshop
- Demonstration of Nagios-based Prototype
- SAM Availability calculation
  - Including equivalence of components across multiple grid infrastructure
- Site Local vs. Central tests
  - What is a good balance?




## Experiment Day

- Morning: experiment plenary presentations
- Afternoon: experiment parallel sessions
- (Detail to be added...)



## Workshop Summary

- There is also a workshop summary talk at the 'beginning of the end' at CHEP...
  - Thursday morning during 'Summary' session...
- Preview key highlights at end of workshop...
- BOFs and parallel session conveners invited to provide 1-2 slide of key results / messages!





## Targetted Interventions

- Common interventions include:
    - Adding additional resources to an existing service;
    - Replacing hardware used by an existing service;
    - Operating system / middleware upgrade / patch;
    - Similar operations on DB backend (where applicable).
  - Pathological cases include:
    - Massive machine room reconfigurations, as was performed at CERN (and elsewhere) to prepare for LHC;
    - Wide-spread power or cooling problems;
    - Major network problems, such as DNS / router / switch problems.
- **Pathological cases clearly need to be addressed too!**



## More Transparent Interventions

- *I am preparing to restart our SRM server here at IN2P3-CC so I have closed the IN2P3 channel on prod-fts-ws in order to drain current transfer queues.*
- *I will open them in 1 hour or 2.*
- Is this a transparent intervention or an unscheduled one?
- A: technically unscheduled, since it's SRM downtime.
- ☺ An EGEE broadcast was made, but this is just an example...
- But if the channel was first paused - which would mean that no files will fail - it becomes instead **transparent** - at least to the FTS - which is explicitly listed as a separate service in the WLCG MoU, both for T0 & T1!
- i.e. if we can trivially limit the impact of an intervention, we **should** (c.f. WLCG MoU services at Tier0/Tier1s/Tier2s)





## Service Review

- For each service need current status of:
  - Power supply (redundant including power feed? Critical? Why?)
  - Servers (single or multiple? DNS load-balanced? HA Linux? RAC? Other?)
  - Network (are servers connected to separate network switches?)
  - Middleware? (can middleware transparently handle loss of one of more servers?)
  - Impact (what is the impact on other services and / or users of a loss / degradation of service?)
  - Quiesce / recovery (can the service be cleanly paused? Is there built-in recovery? (e.g. buffers) What length of interruption?)
  - Tested (have interventions been made transparently using the above features?)
  - Documented (operations procedures, service information)