

# CCRC'08 AND ALL THAT

## Status & Outlook for 2008

[Jamie.Shiers@cern.ch](mailto:Jamie.Shiers@cern.ch)

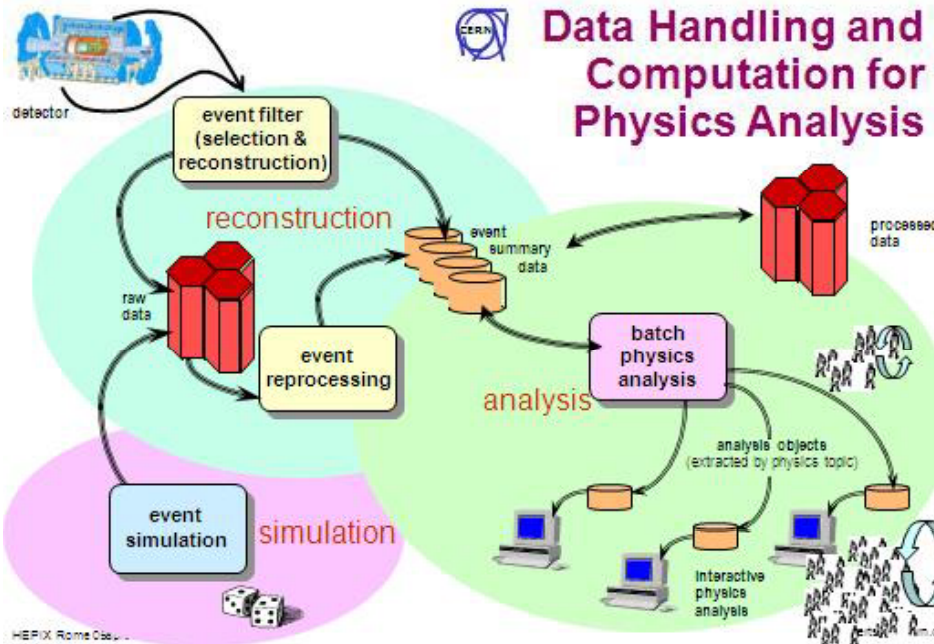
9<sup>th</sup> January 2008

# Introduction

- 2008 promises to be a very busy – and hopefully very rewarding year
- This is the year in which collisions in the LHC are planned and will exercise all aspects of the WLCG Computing Service at all sites and for all supported experiments concurrently
- **The goal of the CCRC'08 exercises is to understand where we stand with respect to these needs and to identify and fix any problems as rapidly as possible**
- **If there are 'no surprises' that will be a surprise in itself!**
- We must **assume** that not everything will work as expected / hoped and work as efficiently as possible towards **solutions**
- The (very) recent past has numerous examples where 'show-stoppers' "disappeared" almost overnight!

# LHC Computing is Complicated!

- Despite high-level diagrams (next), the Computing TDRs and other very valuable documents, it is very hard to maintain a high-level view of all of the processes that form part of even one experiment's production chain
  - See also "First 3 days in the life of a CMS event" (URL in notes)
- Both the detailed views of the individual services, together with the high-level "WLCG" view are required...
- It is ~impossible (for an individual) to focus on both...
- Need to work together as a team, sharing the necessary information, aggregating as required etc.
- **The needed information must be logged & accessible!**
- (Service interventions, changes etc.)



## LCG Service Hierarchy

### Tier-0 - the accelerator centre

- Data acquisition & initial processing
- Long-term data curation
- > Data Distribution to Tier-1 centres



### Tier-1 - "online" to the data acquisition process

- high availability
- Managed Mass Storage - → grid-enabled data service
- > All re-processing passes
- Data-heavy analysis
- National, regional support

### Tier-2 - ~100 centres in ~40 countries

- Simulation
- End-user analysis - batch and interactive
- > Services, including Data Archive and Delivery, from Tier-1s

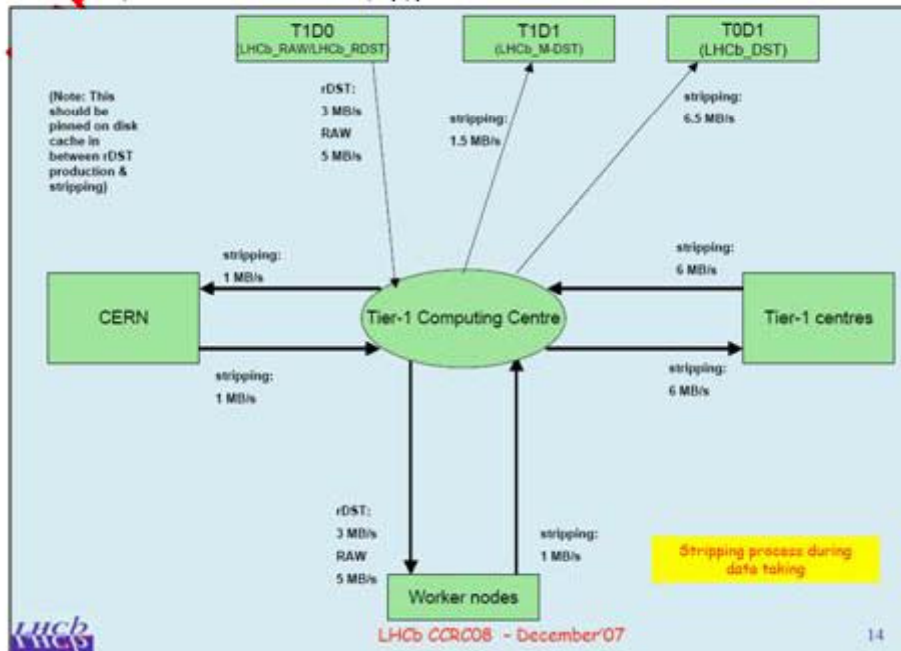
HEPIX Rome Obs...

HEPIX Rome Obs...

les.robertson@cern.ch

LHCb (Nick Brook) Presentation at CCRC08 meeting 4-Dec-2007:

<http://indico.cern.ch/materialDisplay.py?contribId=5&materialId=slides&confId=20248>



WLCG - 26 November 2007

## High: Tier-0 Operations

Tier	Service	Criticality	Consequences of service interruption
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	Online-offline database connectivity	High	No export of conditions data, this delays offline calibration and reconstruction. Import of calibration and configuration data proceed via files and does not require this connection.
0	Castor internal data movement	High	Slow down or interruption of Tier-0 processing. Buffers will fill up after 5 days.
0	Tier-0 processing farm	High	Slow down or interruption of Tier-0 processing. Buffers will fill up after 5 days.
0	Oracle database RAC (offline, ATLR)	High	Slow down or interruption of Tier-0 processing. Buffers will fill up after 5 days. No export of database data.


Dario Barberis: ATLAS Critical Services

## A Comparison with LEP...

- In January 1989, we were expecting  $e^+e^-$  collisions in the summer of that year...
- The “MUSCLE” report was 1 year old and “Computing at CERN in the 1990s” was yet to be published (July 1989)
- **It took quite some time for the offline environment (CERNLIB+experiment s/w) to reach maturity**
- Some key components had not even been designed!
- **Major changes in the computing environment were about to strike!**
- We had just migrated *to* CERNVM – the Web was around the corner, as was distributed computing (SHIFT)
- (Not to mention OO & early LHC computing!)



# Startup woes – BaBar experience



# WLCG “Calendar”

- In October, we came up with an outline schedule for the first half of this year
- Some attempt to turn this into milestones
- **Almost certainly need to maintain a “WLCG Calendar” – which is bound to change...**
- Review ~monthly, e.g. at GDBs
- Like a weather forecast, accuracy will decrease as one looks further & further ahead...
- **But will be essential for planning the service – e.g. migrations of DB services to new h/w at Tier0 – as well as vacations!**

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	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	
<h2>More detail needed...</h2>				
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		Easter 21-24 March
Apr	ALICE ATLAS CMS LHCb	FDR phases 1-3		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



# Reviewing Progress...

- Quarterly reports; LHCC referees reviews
  - Weekly reports to MB; monthly reports to GDB
  - Workshops planned for April (21-25) as well as June
  - ¿ **Should the latter remain at CERN (June 12-13) or move with GDB to Barcelona?**
  - In depth technical analysis needs more & different people to those who attend GDBs...
  - The above is all in place and well understood / exercised
- **Shorter term (daily) follow-up also required**
- ➔ **WLCG Service Coordination role (more tomorrow...)**

# WLCG / EGEE / EGI Timeline

- In 2010, the LHC will reach design luminosity
- In 2010, EGEE III will terminate
- It is inconceivable that we:
  - a. Don't run the LHC machine
  - b. Run the LHC machine without a computing infrastructure (Grid)
  - c. Run the computing infrastructure without Grid operations
- **This is required for other mission critical applications that are dependant on this infrastructure**
- The transition to the new scenario must be
  - a. On time
  - b. Non-disruptive
- **This is a fundamental requirement – it is not an issue for discussion (and is one of EGI\_DS design principles)**

# EGI\_DS: WP3 Milestones & Deliverables

- M3.1 Presentation of First Schema of EGI functions, options analysis and the draft Convention to the NGIs (month 7)
- D3.1: First EGI Functions Definition—Functions, success models, relationship between EGI and NGI, need for new projects (month 9)
  - D3.1.1 – Survey of European & National projects
  - D3.1.2 – Handover from WP2 & WP6
  - D3.1.3 – First Schema of EGI Functions
- D3.2: Final EGI Functions Definition (month 15)
- **WP3 workshop at CERN Jan 29-31 which will address D3.1.3**
  - **WP3 members + invited experts**
- **Proposal: devote 1 day during April Collaboration workshop to review operations “best practices” – both to assist EGI & for our own benefit in the medium – long term**

# EGI\_DS: WP3 Milestones & Deliverables

- M3.1 Presentation of First Schema of EGI functions, options analysis and the draft Convention to the NGIs (month 7)
  - D3.1: First EGI Functions Definition—Functions, success models, relationship between EGI and NGI, need for new projects (month 9)
    - ✓ **D3.1.1 – Survey of European & National projects**
    - + D3.1.2 – Handover from **WP2** & WP6
    - ✓ **D3.1.3 – First Schema of EGI Functions**
  - D3.2: Final EGI Functions Definition (month 15)
  - **WP3 workshop at CERN Jan 29-31 which will address D3.1.3**
    - **WP3 members + invited experts**
- **Proposal: devote 1 day during April Collaboration workshop to review operations “best practices” – both to assist EGI & for our own benefit in the medium – long term**



# Reliable Services – The Techniques

- ☺ DNS load balancing
- ☺ Oracle “Real Application Clusters” & DataGuard
  - H/A Linux (less recommended... because its not really H/A...)
  - 💣 **Murphy’s law of Grid Computing!**
  - Standard operations procedures:
    - Contact name(s); basic monitoring & alarms; procedures; hardware matching requirements;
  - **No free lunch! Work must be done right from the start (design) through to operations (much harder to retrofit...)**
  - Reliable services take less effort(!) to run than unreliable ones!
  - 💣 **At least one WLCG service (VOMS) middleware does not currently meet stated service availability requirements**
  - 💣 **Also, ‘flexibility’ not needed by this community has sometimes led to excessive complexity (complexity is the enemy of reliability) (WMS)**
  - **Need also to work through experiment services using a ‘service dashboard’ as was done for WLCG services (see draft [service map](#))**



# Middleware Developers' Tips

- *"The key point about designing middleware for robustness and resilience is to incorporate these aspects into the initial design.*
- *This is because many of the deployment and operational features already discussed have an impact on the basic architecture and design of the software; it is typically much more expensive to retrofit high-availability features onto a software product after the design and implementation (although it is possible).*
- *In general, the service scaling and high-availability needs typically mandate a more decoupled architecture."*
- **See presentation for more details plus paper for EELA 3**
- **(It's hard to summarize a workshop whilst its going on - let alone write 3 + 1 paper!)**



# WLCG Service Reliability - Goals

- Measured improvements in service quality: April workshop
- Monitor progress using a 'Service Map'
- Size of box = criticality; colour = status wrt "checklist"
- CHEP 2009: all main services at required service level
- Database(-dependent) and data / storage management services appear (naturally) very high in the list + experiment services!
- 24x7 stand-by rota should be put in place at CERN for these services, at least for initial running of the LHC



# Service Reliability: Follow-up Actions

1. Check m/w (prioritized) against techniques - which can / do use them and which cannot? → Priorities for development (service)
2. Experiments' Lists of critical services: service map (FIO+GD criteria)
3. Measured improvement - how do we do it?
4. VO Boxes → VO services
5. Tests - do they exist for all the requested 'services'? → SAM tests for experiments
6. ATLAS & CMS: warrant a dedicated coordinator on both sides
7. Database services: IT & experiment specific
8. Storage - does this warrant a dedicated coordinator? Follow-up by implementation
9. Revisit for Tier1s (and larger Tier2s)
- 10. Overall coordination? → LCG SCM → GDB → MB/OB**
11. Day 1 of WLCG Collaboration workshop in April (21<sup>st</sup>)
12. Long-term follow-up? → solved problem by CHEP 2009
13. "Cook-book" - the current "knowledge" is scattered over a number of papers - should we put it all together in one place? (Probably a paper of at least 20 pages, but this should not be an issue.)




## VOBOX Hardware:

- Resource requirements and planning
  - it is not always easy to have an additional disk on demand because “/data” becomes full
- Hardware warranty
  - Plan for hardware renewal
  - Check warranty duration before moving to production
- Hardware naming and labeling
  - Make use of aliases to facilitate hardware replacement
  - Have a “good” name on the sticker
    - e.g. All Ixbiiii machines may be switched off by hand in case of a cooling problem

**💣 Some “critical services” run over Xmas were just that – and nodename hard-coded in application!**

Activities	Results
1. Crisis & problems	Stress, burn-out, fire fighting, crisis management
<b>2. Planning, new opportunities</b>	<b>Vision, balance, control, discipline</b>
3. Interruptions, e-mail, ...	Out of control, victimised
4. Trivia, time wasting	Irresponsible, ...



# Summary & Conclusions

- We have a very busy and challenging year ahead of us
- Using the **extensive experience** from previous colliders; the **many** data & service challenges and production activities over the past years we are in **good shape** to address the **remaining** challenges for first LHC data taking
- It won't be easy – it won't always be fun – but it'll work!
- **All key infrastructures / roles need to be in place now & validated – if necessary adjusted – in February run**