

A hallway with large projection screens displaying particle detector data. The screens show various views of the detector, including cross-sections and data plots. The text is overlaid on the image.

# LHC: First Contact


**Experience with first LHC data in the  
Worldwide LHC Computing Grid (WLCG)**

**Jamie.Shiers@cern.ch**

**Grid Support Group, IT Department, CERN**

# Agenda

---




- The world's largest scientific machine is undergoing final preparations for data taking from pp collisions ( $5 + 5$  TeV) later in 2008
- The data collected – some 15PB per year – will be analyzed by collaborations of thousands of scientists from hundreds of sites around the world
- This will be done using a worldwide (virtual) grid – federating resources from several large-scale grid infrastructures – to offer a high (performance, quality) production service




# The LHC

- The LHC is a particle accelerator project of international collaboration
- It is made up of several machines, including the Superconducting Proton Synchrotron (SPS) whose protons are injected into the LHC machine
- First proposed in 1984, it took some 15 years to build
- I personally worked on it since 1992 (The rest worked on the collider in the 1970s)



centric  
of large scale  
enuity  
ologies –  
idity –  
sign of the  
has been  
computing  
t CERN!  
previous

# Motivation





- We currently have a good and very accurate model that has been extensively validated by experiment
- ☹ But it is – at best – incomplete (or possibly **wrong**), leaving some important open questions:
  - Mass;
  - Matter vs anti-matter;
  - Dark Matter;
  - Dark Energy
- The LHC has been built as a **Discovery Machine** to hopefully answer these questions – and perhaps raise some more!



# CERN




- To some people, CERN is simply a geographic location
  - Latitude: 46°13'59" N
  - Longitude: 6°3'20" E
- For me this description is more than incomplete – it is simply *wrong!*
- I do not believe that you can really understand what CERN is (about) unless you also consider:
  - The scientific research programme;
  - The close collaboration with a large number of institutes worldwide – that CERN serves and for whom it exists;
  - Its outreach programme and technology transfer; ...
- We need to recognise (explain, evangelise) the role of science in society – it is not just for science!



- For the purpose of this talk I will use a similarly loose definition of the **Worldwide LHC Computing Grid**
- This formally consists of a collaboration between the 4 main LHC “experiments” and a set of institutes that provide computing resources and services to these communities
  - Defined in a “Memorandum of Understanding” signed by all parties;
  - Includes services and service levels offered, resource pledges for coming years
- Also a set of management infrastructures and operational boards to plan, deploy, operate and evolve the services (**close collaboration with EGEE etc.**)
- IMHO, essential to also include “friends” in this informal definition – other (VOs, sites, services) with various couplings to the “core business”
  - e.g. GEANT4 – main simulation tool in HEP and (way) beyond...
  - [SIXT](#) – simulation tool for LHC accelerator (sixtrack)
  - Lattice QCD simulation – a number of serious scientific publications (see slide notes)
- Many physicists / institutes also involved in other experiments and / or disciplines
- The “boundaries” – if they even exist – are tenuous & flexible...
  - e.g. CERN Grid Support group works with a wide range of disciplines;
  - Grid Data & Storage Management products (dCache, DPM, ...) – even if developed at HEP labs for HEP users are also used by – and extended for – many disciplines...

# Data Handling and Computation for Physics Analysis





# LHC: One Ring to Bind them...

## Introduction

Status of

LHCb

ATLAS

ALICE


CMS

Conclusions



LHC : 27 km long  
100m underground



pp, B-Physics,  
CP Violation



General Purpose,  
pp, heavy ions



Heavy ions, pp








# ATLAS Cloud Model

## Tier-0 - the accelerator centre

- Data acquisition & initial processing
- Long-term data storage
- **Data Distribution**




## Tier-2 - "regional support" centres

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



# Tier 0 at CERN: Acquisition, First pass reconstruction, Storage & Distribution





# Tier 0 – Tier 1 – Tier 2

## Tier-0 (CERN):

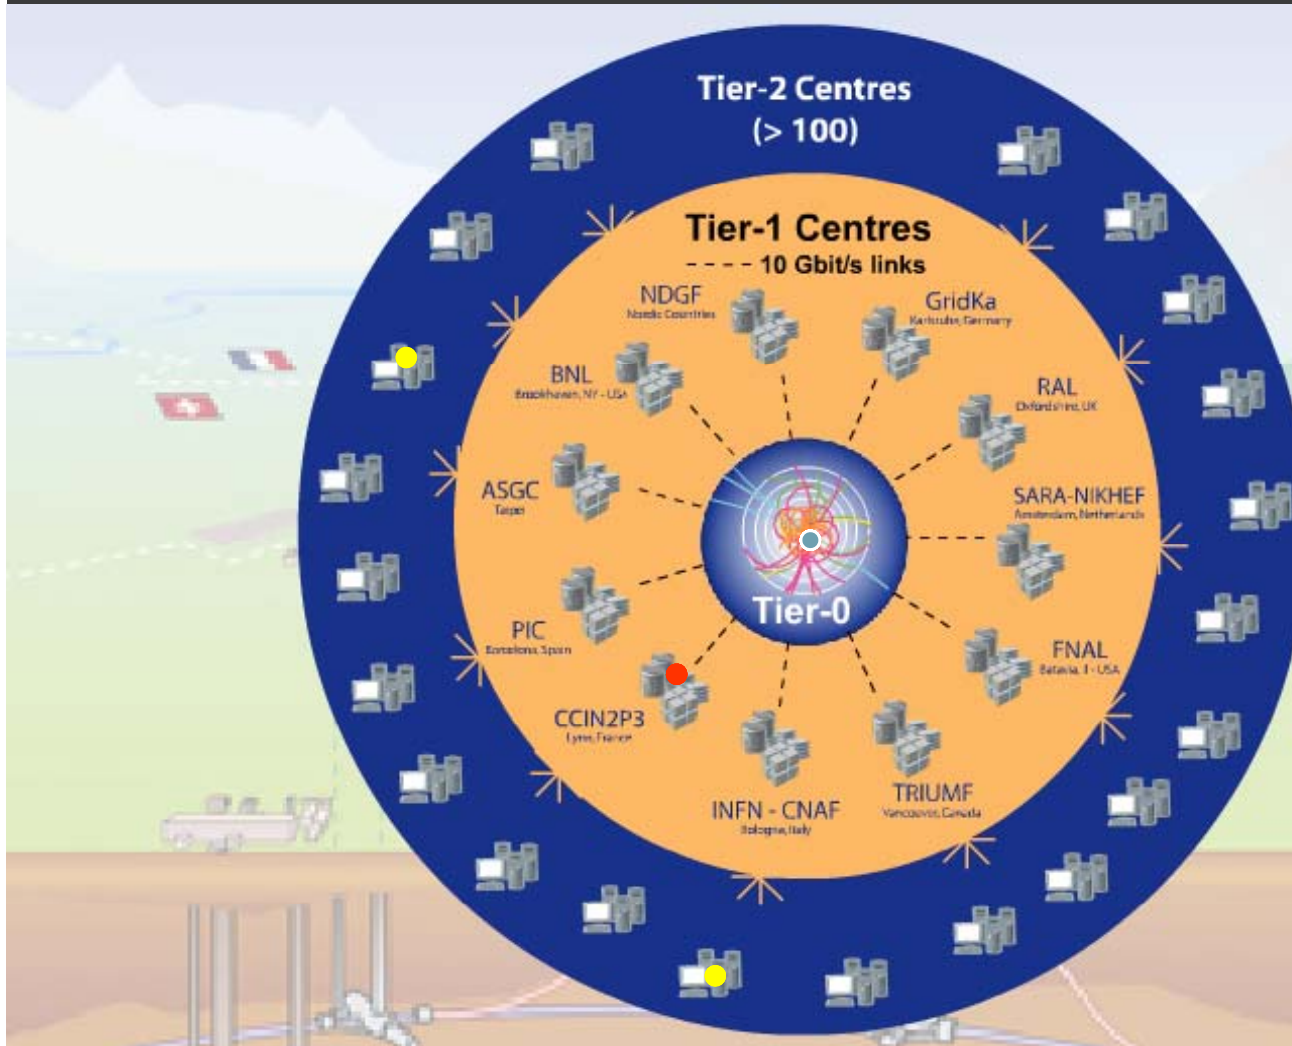
- Data recording
- Initial data reconstruction
- Data distribution

## Tier-1 (11 centres):


- Permanent storage
- Re-processing
- Analysis

## Tier-2 (>200 centres):

- Simulation
- End-user analysis









BNL



ASGC/Taipei



CCIN2P3/Lyon



TRIUMF/BC



NIKHEF/  
SARA



FNAL



RAL



PIC



NDGF



CNAF



TIER2s



CERN



FZK



# WLCG Service



How has the service stood up to real production usage?



# Service Status – The Story So Far...

---

- One year ago we had still not demonstrated that we could sustain all production workflows from all 4 LHC experiments simultaneously
- This prompted the “legendary question”:
  - 💣 **What happens when the LHC is operating?**
- This led to the “Common Computing Readiness Challenge(s)” that were exercised during the first half of this year
  - Agreed metrics, targets & reporting mechanisms...
- The conclusion from the challenge (February and May) was that:
  - We met the goals (even if overlap from all experiments less than optimal) **but**
  - **Real data taking will be different!**
- The real – and very frightening – prospect had CCRC’08 been less successful would have been de-scoping!
- ☺ This option was ruled out already by the February run
- **IMHO – the success of CCRC’08 is a landmark in the fable of grid computing (and obviously to the many people who contributed to this) and deserves(d) much more than a couple of bullets in passing...**

## How We Measured Our Success

- Agreed up-front on specific targets and metrics – these were 3-fold and helped integrate different aspects of the service

Experience shows this is not enough! Computing models of experiments must also be considered.

2. ... for each site to ensure ... in place;
3. WLCG **“Memorandum of Understanding”** (MoU) targets – services to be provided by sites, target availability, time to intervene / resolve problems ...

↳ **Clearly some rationalization of these would be useful – significant but not complete overlap**



## Problem Response Time and Availability targets Tier-1 Centres

<i>Service</i>	<i>Maximum delay in responding to operational problems (hours)</i>			<i>Availability</i>
	<i>Service interruption</i>	<i>Degradation of the service</i>		
		<i>&gt; 50%</i>	<i>&gt; 20%</i>	
<b>Acceptance of data from the Tier-0 Centre during accelerator operation</b>	<b>12</b>	<b>12</b>	<b>24</b>	<b>99%</b>
<b>Other essential services – prime service hours</b>	<b>2</b>	<b>2</b>	<b>4</b>	<b>98%</b>
<b>Other essential services – outside prime service hours</b>	<b>24</b>	<b>48</b>	<b>48</b>	<b>97%</b>



# Critical Service Follow-up

- Targets (not commitments) proposed for Tier0 services
  - Similar targets requested for Tier1s/Tier2s
  - Experience from first week of CCRC'08 suggests targets for **problem resolution** should not be too high (if ~achievable)
    - The MoU lists targets for responding to problems (12 hours for T1s)
- ¿ Tier1s: 95% of problems resolved < 1 working day ?
- ¿ Tier2s: 90% of problems resolved < 1 working day ?
- **Post-mortem triggered when targets not met!**

Time Interval	Issue (Tier0 Services)	Target
End 2008	Consistent use of all WLCG Service Standards	100%
30'	Operator response to alarm / call to x5011 / <b>alarm e-mail</b>	99%
1 hour	Operator response to alarm / call to x5011 / <b>alarm e-mail</b>	100%
4 hours	Expert intervention in response to above	95%
8 hours	Problem resolved	90%
24 hours	Problem resolved	99%

# CCRC'08 Post-Mortem High-lights

---

- The bottom line – we believe that the experience in 2008 so far confirms that we have a working service model and that we are ready to **face** the challenges of data taking from pp collisions in the LHC
  - ✓ Most aspects of the service work well most of the time
  - ✓ We have a proven track record in resolving even the most daunting of problem in an acceptably short time
- What is really interesting is what happens when things go wrong – and how we can improve on this in the future

# Strengths

- CCRC'08 and accompanying experiment “dress rehearsals” have in most cases demonstrated that the services / sites / experiments are ready for higher loads than are expected from 2008 pp data taking
- ☺ **The middleware process is working well!**
- ☺ **The database services are working well!**
- **We have a well tested service model and have demonstrated steady improvement over a long time**

# Weaknesses

- Some of the services – including but not limited to storage / data management – are still not sufficiently robust. (Process? Deployment?) We have (so far) failed to define and regularly update a clear table of versions + release + patch level. This is nevertheless the target, with a weekly update at the joint EGEE-OSG-WLCG operations meeting
- Communication is still an issue / concern. This requires work / attention from everybody – it is not a one-way flow.
- Not all activities (e.g. reprocessing, chaotic end-user analysis) were fully demonstrated even in May, nor was there sufficient overlap between all experiments (and all activities). Work continues (July and beyond)...
- There were a large number (IHMO too many) Tier0 service upgrades in June – not always well scheduled and / or motivated. We must balance stability with needed fixes



# Opportunities

- There is no technical reason why we cannot solve the non-technical problems in the storage area (i.e. define recommended versions that have been released and tested – not “dreams”!)
- Communication – certainly no silver bullet expected. Need solutions that scale to the number of sites / players involved, that can adapt to constraints of time zones and affordable technology (audio & video conferencing, for example...)
- Improvements in monitoring and automation to reduce human expert involvement to a sustainable level (medium – long-term?)
- We still need to maintain a high(-er) level view of the overall WLCG service – a purely component view is not compatible with a highly complex service with many inter-dependencies

# Threats

- The biggest threat that I see is to fall back from reliable service mode into “fire-fighting” at the first sign of (major?) problems.
- This in the past has been accompanied by memos to the highest level, triggering time and effort consuming response / post-mortems, but is not sustainable and is much less efficient than the proven service mode.
- This requires close collaboration and concerted effort – as has been the case through many years of data and service challenges, and as we have seen at previous machines.
- Daily operations meeting as a focus / dispatching point plus constant interactions with experiments / sites.

# S.W.O.T. Summary

- CCRC'08 has proven to be a very valuable exercise for demonstrating readiness for 2008 data taking, including identifying (and fixing) holes in the service
- With justification, we can be confident of our readiness – from steady operation through to unexpected “crises” (which we will quickly defuse & resolve...)
- Communication & coordination have been key
- It has been – at least at times – very hard work, but also extremely rewarding!
- May collisions commence...



# Post-CCRC'08

---

- Since the end of the formal challenge, some of the outstanding workflows – e.g. re-reconstruction at Tier1s including conditions DB access – have been successfully demonstrated
- Many services (-ers) have been upgraded!
- The “summer effect” (numerous absences, sometimes long delays in problem resolution) was as marked as in previous years...
- 💣 **Had we been taking data this summer it would have been somewhat painful, with prolonged service and even site downtime(s)**
- This leads to two major issues that are still not satisfactorily resolved:
  - How to plan, schedule and perform (disruptive) interventions – also taking into account multi-site issues (i.e. not all Tier1s down together please...)
  - How to provide adequate service coverage (“summer” being a particular, but by no means the only, case).

# Storage Issues (and beyond...)

---

- Configuration still an issue
  - **Conceptually, this does not seem so complicated.**
  - But in reality, it has taken a lot of time and many iterations to
    - Understand what is possible;
    - Map this to what is required;
    - Describe in a clear and unambiguous way what is required by VO and by site;
    - Get it implemented.
- Release Procedures
  - Bug fix releases have in the very recent past included things marked as “changed” “added” or “removed”.

# WLCG Service



Beyond the Standard Model



# Operations Introduction

---

- ❑ I am not going to describe the standard EGEE operations, application and user support model(s)
- ❑ I will mention a few of the things (procedures, tools) that have been essential during the WLCG Service Challenges and CCRC'08 – these are expected to continue at least for the immediate term
- ❑ I will also describe “WLCG Operations” – again taking a very broad brush, covering not only the “WLCG Service” but also that part of the experiment-specific operations that is exposed to it
- ❑ **The bottom line – for large-scale production, there is a (very) significant amount of effort on top of what is provided by the basic infrastructures**

Find

- All CERN
- IT Department
- Local Site

## IT Services

[Access to Facilities](#)  
[Desktops & Portables](#)  
[Physics Computing](#)  
[Technical Computing](#)  
[Telecom](#)  
[Index of Services](#)  
[Other Services](#)

## Divisional Information

[What We Do](#)  
[Divisional Structure](#)  
[Divisional Administration](#)  
[Minutes & Reports](#)  
[Safety in IT Division](#)  
[Job Opportunities](#)

## More on Computing

[CERN School of Computing](#)  
[Colloquia](#)  
[Seminars](#)  
[Newsletter](#)  
[Bookshop](#)

# IT Service Status Board

(See also the [Detailed Services Status](#))

Given the increased publicity related to the LHC start-up vigilance is needed to protect against a possible increase in attempts to break-in to CERN computing infrastructure. Please use your computer account(s) with care and vigilance, in particular when browsing external web pages.

**Do not reply to doubtful emails, do not click on links within doubtful emails, and delete suspicious email attachments without opening them. Never respond to a request to give your username and password, credit card details, or any similar information as such requests are always bogus.**

([Read more...](#))

Date	Description	Affecting, explanation	Posted / Comment
17 September Time 10:20	<a href="#">Problem with Network Connection Request Form</a>	<a href="http://network.cern.ch">http://network.cern.ch</a> was unaccessible for a short period (apparently between 10:20 and 10:27).	17-09-08
16 September Time: 13:52	<a href="#">Emergency Power Cut in Preveessin at 13:52</a>  <b>Update at 16:45</b>	An emergency power cut has taken place in Preveessin Building 867, which triggered many failures of the computing network. The affected buildings are: <b>867, 974, 904, 927, 892, 926</b> <b>Most problems fixed, except in the ground floor of building 904</b>	16-09-08
10 September (at 10:13)	<a href="#">CMSMON web page defaced - Security issues in Web applications</a>	Security issues in Web applications are a common threat, and the CERN Security Team recommends all service managers to review the security of their web applications. More information is available at: <a href="http://cern.ch/security/webapps/">http://cern.ch/security/webapps/</a>	15-09-08  Further details, see link on the left
As from 12 September (8:00)	<a href="#">Problem with printer 513-HP-A0GLOSSY in building 513</a>	The printer 513-HP-A0GLOSSY located in building 513 and used to produce posters on glossy paper is currently down and will not be repaired until next week. For urgent requests, please contact the Computing Helpdesk at 78888 or <a href="mailto:Computing.Helpdesk@cern.ch">Computing.Helpdesk@cern.ch</a>	12-09-08  Sorry for the inconvenience.

[Archive](#)

## Scheduled Interventions

Scheduled on	Intervention	Reason/Affecting	Posted / Comment
15 September  Update forced on Tuesday 16 September	<a href="#">NICE Updates for IT Computers</a>	Important updates will be applied to all NICE workstations within IT Department. A reboot is required after the installation. The update will be forced on Tuesday, 16th September. You can find all relevant details <a href="#">here</a>	12-09-08
18 September	<a href="#">Short intervention on</a>	This intervention will be short (~15 minutes) but some	17-09-08

Item

Target

0%

%

0%

%

%

%

# WLCG Operations

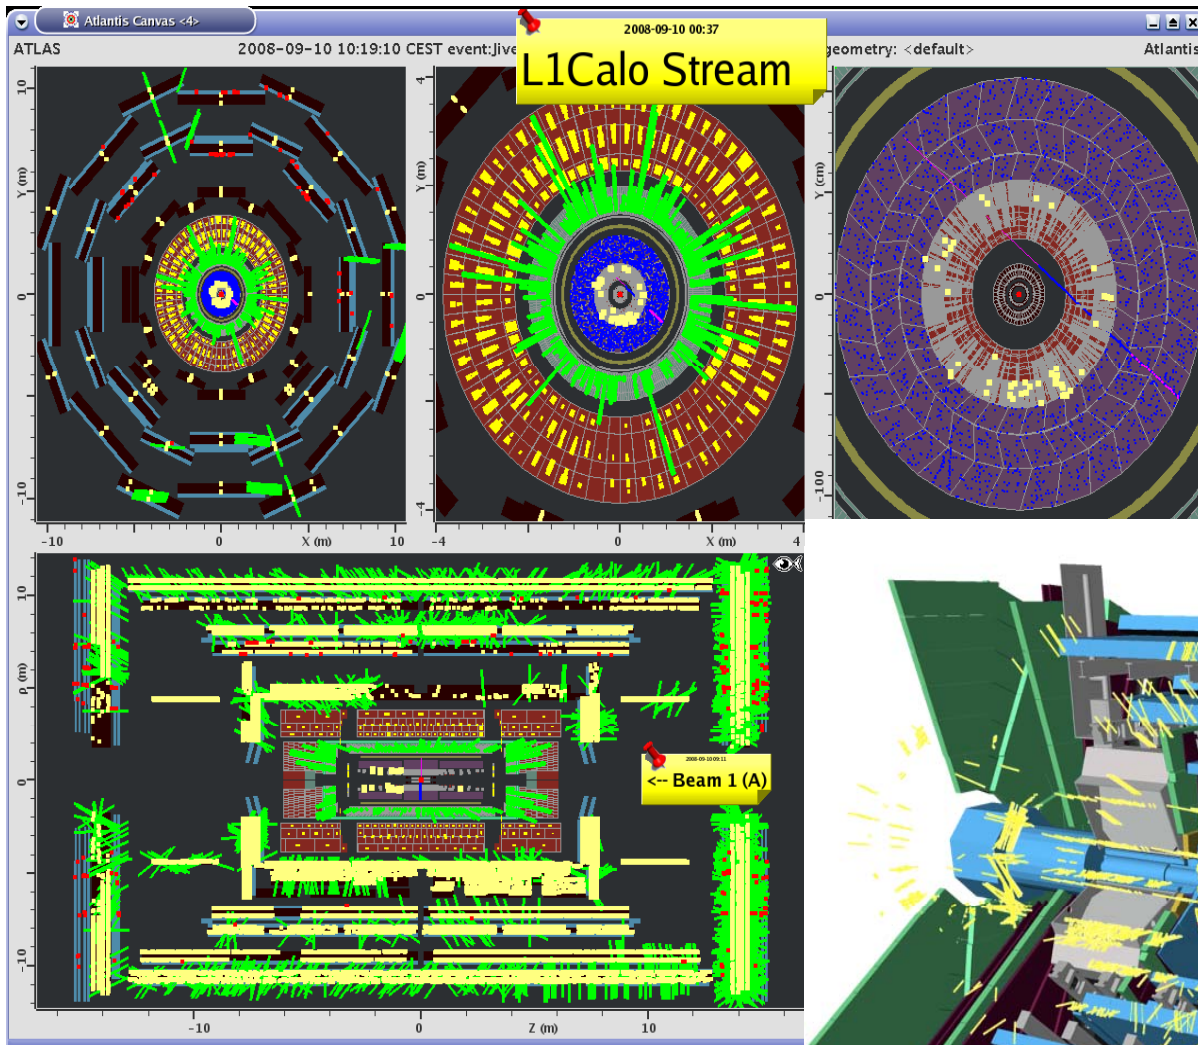
---

- Whilst it is understood that this builds extensively on the operations infrastructure(s) of the underlying grids, there are additional layers that have proven at least valuable (so far...)
- These include:
  - Daily operations con-calls; week-days at 15:00 Geneva time, notes distributed same business day and widely read by Tier1 (Tier2?) sites, experiments and WLCG management
  - Weekly service summary to WLCG Management Board, quarterly service summary to Overview Board
  - Additional follow-up of service problems (beyond SAM service availability at MB): service issues that “violate” MoU target(s) trigger a post-mortem, which should normally be available by time of following MB
- The experiments also have extensive operations teams & infrastructures, e.g.
  - WLCG Collaboration workshops: 200-300 people;
    - **Hold jointly with other events, e.g. CHEP, EGEE'09(?), where possible**
  - ATLAS “jamborees”: closer to 100...




***... as well as in the ATLAS Tier-0 Control Room:  
Analysing the first LHC events on 10<sup>th</sup> September 2008***





Online display

Offline display



first beam event seen in ATLAS

The very first beam-splash event from the LHC in ATLAS on 10:19, 10<sup>th</sup> September 2008



# WLCG VO Operations

---

- The Experiment Dashboards, VO-specific SAM tests, together with other experiment-specific monitoring really are used to “run the show”
  - CMS: 15’ daily to check status of all sites!
  - IMHO – no point to attempt to optimize this (already great) result just yet – get more experience with real data taking!
- **Very close collaboration between Grid Support team and experiments / sites**
- **Very close collaboration with Grid & Service experts at CERN and elsewhere**
- Compares favorably with other models – you really need experts both in relevant technology as well as computing model of application in question to provide support efficiently & effectively!

# WLCG Service



Immediate Concerns and  
Challenges

# My top three concerns...

---

1. Dealing with the **realities** of data taking and production which are **bound** to be different to planned exercises including **many more users** and the **pressures** of getting some results;
2. Handling **larger** changes such as new **architectures** that are bound to come (or are already here);
3. **Manpower** and **funding** issues in the post-EGEE III era.



# CCRC '09 - Outlook

---

- SL(C)5
- CREAM
- Oracle 11g
- SRM v2.2++
- ☛ **Other DM fixes...**
- SCAS
- [ new authorization framework ]
- ...
- 2009 resources
- 2009 hardware
- Revisions to Computing Models
- EGEE III transitioning to more distributed operations
- Continued commissioning, 7+7 TeV, transitioning to normal(?) data-taking (albeit low luminosity?)
- New DG, ...

# But what of Grids vs Clouds?

---

- ❑ An issue that was already a concern at the time of LEP was the effective “brain-drain” from collaborating institutes to CERN or other host laboratories
- ❑ Grid computing has (eventually) shown that participating institutes can really make a major contribution (scientifically and through computing resources)
- ❑ This gives valuable feedback to the funding bodies – one of the arguments against “CERNtralizing” everything
- ⚠ ***Can these arguments apply to cloud computing?***
- ❑ Not (obviously) with today’s commercial offerings...
- ❑ But maybe we are simply talking about a different level of abstraction – and moving from today’s complex Grid environment to really much simpler ones
  - Which would be good!
- ❑ But again, what large scale database / data management application really works without knowing (and controlling) in detail the implementation!

# Conclusions

---

- Compared to the criteria in Ian Foster's "What is a Grid? A 3-point checklist" WLCG gets **full marks!**
- The concepts of **collaboration** and **community** have been **essential** in achieving this success – over and beyond the technical successes of the underlying infrastructures
- Whilst we have achieved a great deal, the **challenges** that lie ahead are significant, important and by no means VO-specific
- **Thanks to all who have made this possible...**

# Grid Computing in 3 Easy Steps

- Today there are many definitions of *Grid computing*:
- The definitive definition of a Grid is provided by [\[1\]](#) Ian Foster in his article "[What is the Grid? A Three Point Checklist](#)" [\[2\]](#).
- The three points of this checklist are:
  1. Computing resources are not administered centrally;
  2. Open standards are used;
  3. Non-trivial quality of service is achieved.

