



# WLCG Status Report

22<sup>nd</sup> March 2009  
Collaboration Board

Ian Bird  
LCG Project Leader






# Last year: Highlights

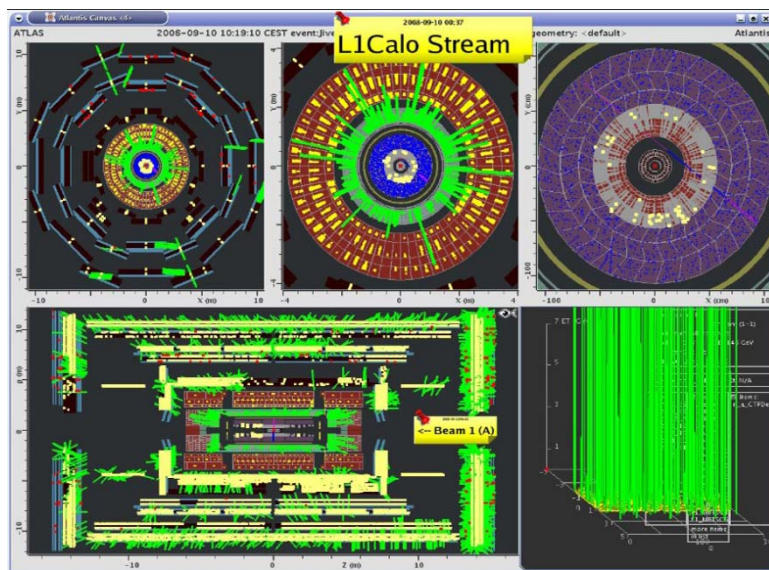

- CCRC'08 was successful
  - But did not fully test everything ...
- Mini reviews in July 2008 and February 2009:
  - “Recommend that there is a CCRC'09 in some form:”
    - At least CMS+ATLAS – but preferably with all 4 experiments
    - Testing reprocessing at Tier 1s (recall from tape) and massive/chaotic user analysis
    - Need metrics with which to evaluate this
  - “Let’s make sure we are not limited by resources when data comes..”
    - Not obvious you can just move the schedule by one year in terms of resources”
  - “Need an official statement on 2009/2010 running time and LHC efficiency factor common for all experiments so they can provide a consistent/coherent estimation of resources needed in 2009/2010”
    - Promised urgently (~today)
  - “Experiments still suffer from SRM (MSS) performance”
  - Applications area – “Very good progress on all fronts with very mature organisation well managed giving results”
- Ready for accelerator start up
  - Took data(!) and cosmics ...
- Mostly the services improved ...




Run 62063, Event 1534



3.1/0.3 fps



Top AfterMagnet Downstream Side






# WLCG MoU Signature Status (Nov C-RRB)




- Since the April 2008 C-RRB the Czech Republic has signed the WLCG MoU as a Tier-2 supporting ALICE and ATLAS
- All Tier-1s have now signed
- All of the Tier-2s have signed except Austria
  - Signature expected before mid November 2008 (done)
- A new MoU will be signed on 11/11/08 with Korea as a Tier-2 supporting CMS (done)
- Brazil is still planning to sign the MoU as a Tier-2 supporting all 4 experiments
  - Letter sent by J. Engelen in June 2008
  - WLCG MoU wording and future commitment being examined by their legal experts


Today:

33 countries have signed (~50 signatures) for 11 Tier 1s,  
61 Tier 2 federations and 120 Tier 2 sites




# CCRC'08 and beyond





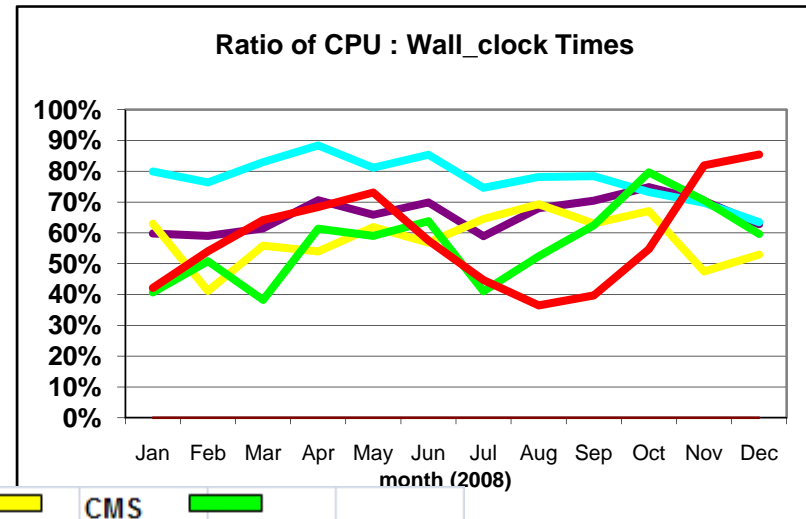

# Usage Patterns

- Can change significantly e.g. between CCRC'08 in May and cosmics/simulations in September
- Tier 2s consistently deliver ~50% of total







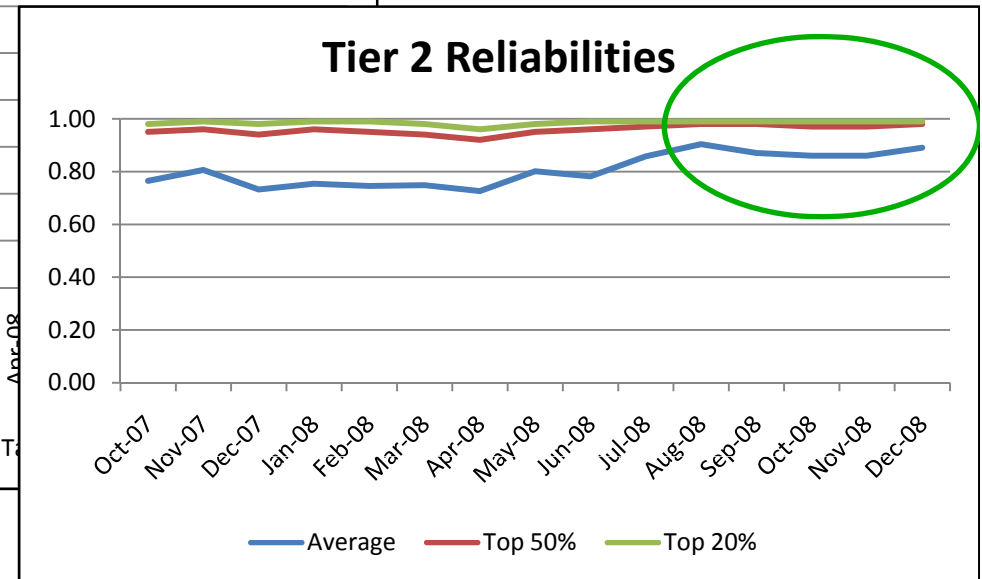

# CERN + Tier 1 accounting - 2008



ALICE █ CMS █  
 ATLAS █ LHCb █  
 installed capacity (inc. efficiency factor) —  
 MoU commitment (inc. efficiency factor) - - -



# Reliabilities



Improvement during CCRC and later is encouraging

- Tests do not show full picture – e.g. Hide experiment-specific issues,
- “OR” of service instances probably too simplistic

- publish VO-specific tests regularly;
- rethink algorithm for combining service instances





# Tier-2 Availability and Reliability Report

Federation Summary - Sorted by Reliability

December 2008

Critical SAM Tests - <http://sam-docs.web.cern.ch/sam-docs/docs/htmldocs/MANUserManual/node22.html>

Availability = % of successful tests  
Reliability = Availability / Scheduled Availability  
Reliab

Colour	Federation	Site	CPU Count	Reliability	Availability	Reliability History			
						Sep-08	Oct-08	Nov-08	
	AT-HEPHY-VIENNA-UIBK ( Austria, Austrian Tier-2 Federation )								
	T2_U	HEPHY-UIBK	262	92 %	92 %	100 %	100 %	99 %	
	T2_U	Hephy-Vienna	324	97 %	97 %	96 %	96 %	95 %	
	FR-C AU-ATLAS ( Australia, University of Melbourne )								
	FI-HI	Australia-ATLAS	80	93 %	93 %	98 %	97 %	94 %	
	US-M	Australia-UNIMELB-LCG2	27	N/A	0 %	3 %	0 %	0 %	
	PT-L BE-TIER2 ( Belgium, Belgian Tier-2 Federation )								
	UK-S	BEgrid-ULB-VUB	332	14 %	12 %	95 %	82 %	98 %	
	CZ-P	BelGrid-UCL	579	96 %	96 %	88 %	97 %	97 %	
	JP-TI CA-EAST-T2 ( Canada-East Federation )								
	RO-L	TORONTO-LCG2	240	41 %	41 %	95 %	97 %	45 %	
	DE-E CA-WEST-T2 ( Canada-West Federation )								
	FR-II	ALBERTA-LCG2	44	98 %	98 %	99 %	98 %	92 %	
	US-M	SFU-LCG2	64	92 %	92 %	87 %	92 %	93 %	
	TW-F	VICTORIA-LCG2	65	95 %	95 %	84 %	93 %	91 %	
	SI-SI			97 %	97 %	97 %	97 %	97 %	
	UK	NorthGrid		97 %	96 %	97 %	97 %	97 %	
	NO	NORDGRID-T2		97 %	97 %	97 %	97 %	97 %	
	DE	DESY-RWTH-CMS-T2		97 %	96 %	97 %	97 %	97 %	
	DE	MCAT		97 %	97 %	97 %	97 %	97 %	
	T2_US	Purdue		95 %	96 %	95 %	95 %	95 %	
	FR	IN2P3-LAPP		96 %	73 %	96 %	96 %	96 %	
	UK	London-Tier2		96 %	95 %	96 %	96 %	96 %	
	FR	IN2P3-IPHC		96 %	95 %	96 %	96 %	96 %	
	HU	HGCC-T2		95 %	95 %	95 %	95 %	95 %	
	AT	HEPHY-VIENNA-UIBK		95 %	95 %	95 %	95 %	95 %	
	CA	WEST-T2		95 %	95 %	95 %	95 %	95 %	
	RU	RDIG		95 %	94 %	95 %	95 %	95 %	
	UK	SouthGrid		95 %	90 %	95 %	95 %	95 %	
	FR	IN2P3-SUBATECH		94 %	94 %	94 %	94 %	94 %	
	US	AGLT2		94 %	94 %	94 %	94 %	94 %	
	IN	DAE-KOLKATA-TIER2		94 %	71 %	94 %	94 %	94 %	
				90 %	90 %	90 %	90 %	90 %	
				85 %	85 %	85 %	85 %	85 %	
				85 %	82 %	85 %	82 %	85 %	
				85 %	85 %	85 %	85 %	85 %	
				84 %	83 %	84 %	83 %	84 %	
				84 %	81 %	84 %	81 %	84 %	
				80 %	77 %	80 %	77 %	80 %	
				80 %	80 %	80 %	80 %	80 %	
				78 %	51 %	78 %	51 %	78 %	
				78 %	79 %	78 %	79 %	78 %	
				66 %	60 %	66 %	60 %	66 %	
				66 %	65 %	66 %	65 %	66 %	
				41 %	41 %	41 %	41 %	41 %	
				0 %	0 %	0 %	0 %	0 %	
				0 %	0 %	0 %	0 %	0 %	
				N/A	N/A	N/A	N/A	N/A	

## Tier 2 reliabilities

Big improvement

Federation average is now weighted by #CPU (where avail)

Would like to fix target at 95%

- Should be achievable

- e.g. of extended scheduled downtimes (availability << reliability)

Only 1 Federation still not reporting (Nordic started in Dec)

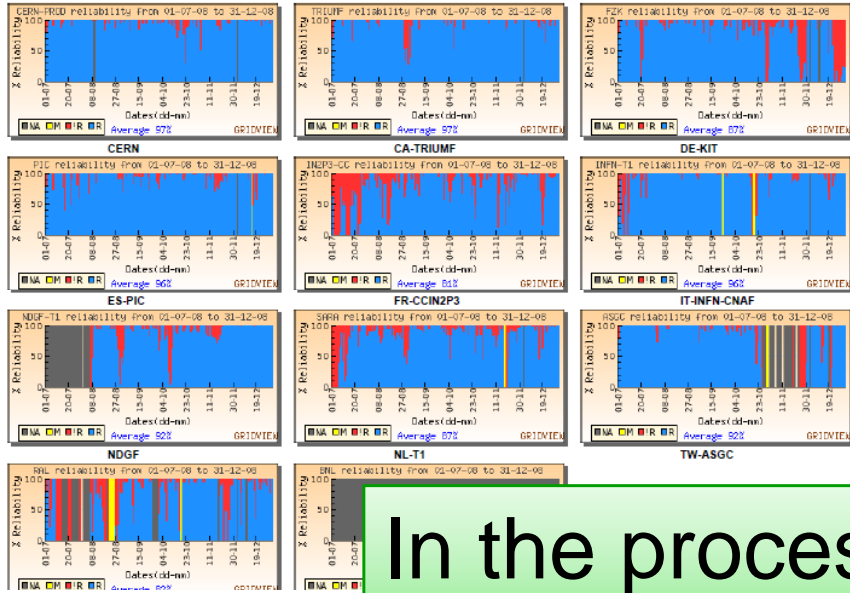
# VO-specific tests



## Reliability of WLCG Tier-1 Sites + CERN for ATLAS

July 2008 - December 2008

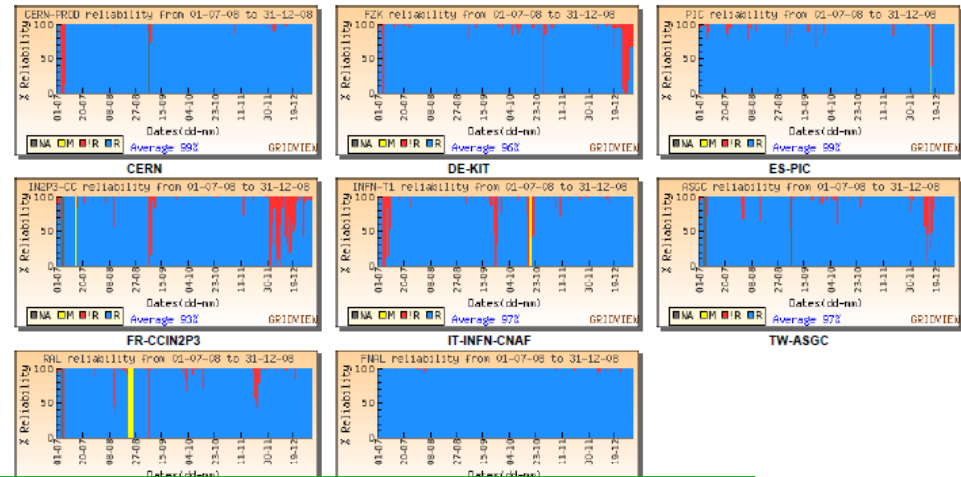
Data from SAM Monitoring. Plots show Reliability for last 6 Months  
Reliability is calculated as  $\text{time\_site\_is\_available} / (\text{total\_time} - \text{time\_site\_is\_scheduled\_down})$   
Target reliability for each site is 95% and Target for 8 best sites is 97% from June, 2008



## Reliability of WLCG Tier-1 Sites + CERN for CMS

July 2008 - December 2008

Data from SAM Monitoring. Plots show Reliability for last 6 Months  
Reliability is calculated as  $\text{time\_site\_is\_available} / (\text{total\_time} - \text{time\_site\_is\_scheduled\_down})$   
Target reliability for each site is 95% and Target for 8 best sites is 97% from June, 2008



In the process of being validated


- December 2008



## Reliability of WLCG Tier-1 Sites + CERN for ALICE

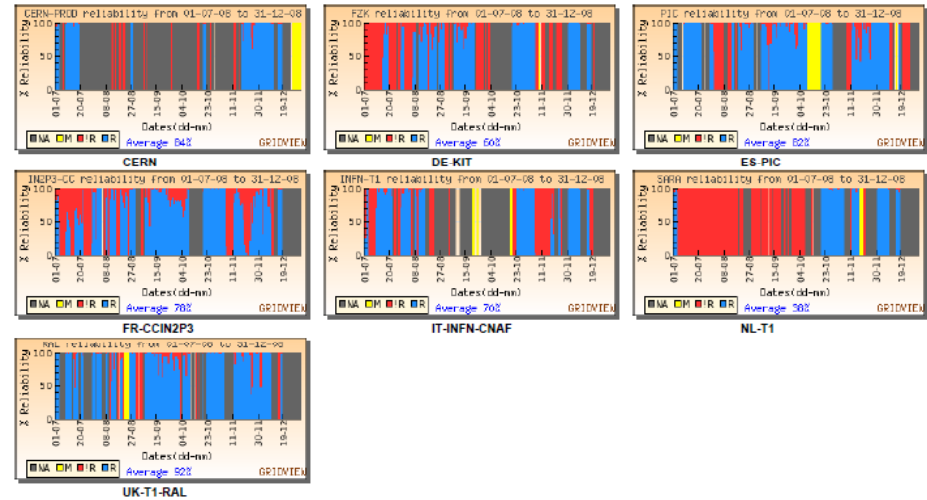
July 2008 - December 2008

Data from SAM Monitoring. Plots show Reliability for last 6 Months  
Reliability is calculated as  $\text{time\_site\_is\_available} / (\text{total\_time} - \text{time\_site\_is\_scheduled\_down})$   
Target reliability for each site is 95% and Target for 8 best sites is 97% from June, 2008



Data from SAM Monitoring. Plots show Reliability for last 6 Months

Reliability is calculated as  $\text{time\_site\_is\_available} / (\text{total\_time} - \text{time\_site\_is\_scheduled\_down})$   
Target reliability for each site is 95% and Target for 8 best sites is 97% from June, 2008






# Pledge Balance in 2009

The table below shows the status at 27/10/08 for 2009 from the responses received from the Tier-1 and Tier-2 sites

- The Total 2009 pledge from Russia is included but not the split across the experiments
- Following a re-organisation of some of the German Federations, pledges for 3 are still to be included
- Pledges for the new French Tier2 IPHC Strasbourg not included

▪ % indicates the balance between offered and required.

	ALICE	ATLAS	CMS	LHCb	Sum 2009
T1 CPU	-49%	6%	-2%	2%	-12%
T1 Disk	-43%	-5%	-13%	-2%	-13%
T1 Tape	-50%	-7%	7%	6%	-13%
T2 CPU	-44%	0%	-8%	-40%	-12%
T2 Disk	-44%	-20%	35%	-	-2%





# Pledge Balance 2008-2013

- The table below shows the global picture for 2008-2013, status as of 17/11/08. % indicates the balance between offered and required
- Some Federations have recently signalled a change to procurements for 2009, not supported by WLCG Management or Overview Boards

	2008	2009	2010	2011	2012	2013
T1 CPU	-5%	-12%	-11%	-15%	-20%	-26%
T1 Disk	-12%	-13%	-15%	-18%	-24%	-29%
T1 Tape	-13%	-13%	-16%	-22%	-24%	-23%
T2 CPU	-2%	-8%	-29%	-31%	-32%	-37%
T2 Disk	-12%	-1%	3%	-6%	-6%	-17%

- Last RRB agreed that planning timescale would change from 5 years to 3 to be more realistic



# Schedule for 2009 - 2010

## Schedule with running in winter months

- Gains 20 weeks of LHC physics (independent of “slip”)

Year	2009												2010														
Month	F	M	A	M	J	J	A	S	O	N	D	J	F	M	A	M	J	J	A	S	O	N	D	J	F	M	
Baseline	SH	SH	SH	SH	SH	SH	SH	SH	SU	PH		SH	SH	SH	SH	SH	SH	SU	PH	PH	PH	PH	SH	SH	SH	SH	
	24 weeks physics possible																										
Base 1	SH	SH	SH	SH	SH	SH	SH	SH	SU	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	SH	SH	SH	SH	
	44 weeks physics possible																										
Gain 20 weeks of physics in 2010 by running during winter months																											
HIGH price Electricity																											
Delay (4W)	SH	SH	SH	SH	SH	SH	SH	SH	SU	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	SH	SH	SH	SH	
Delay (8W)	SH	SH	SH	SH	SH	SH	SH	SH	SH	SU	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	SH	SH	SH	SH	

From Chamonix summary: <http://indico.cern.ch/conferenceDisplay.py?confId=45433>



# Likely scenario

- Injection: end September 2009
- Collisions: end October 2009
- Long run from ~November 2009 for up to 44 weeks
  - Short stop (2 weeks) over Christmas/New Year
- Energy will be limited to 5 TeV
- Heavy Ion run at the end of 2010
  - No detailed planning yet
  
- Now understand the effective amount of data taking in 2009+2010 will be  $\sim 6.1 \times 10^6$  seconds (cf  $2 \times 10^7$  anticipated in original planning)
- Experiments now re-assessing their requirements
  - However, must ensure that there are sufficient resources to rapidly exploit the data
- New requirements discussed on March 31
- Agreement on April 7 of what is to be presented to RRB at end of April





# Resources ...

- **New benchmark agreed**
  - kSI2K → HEP-SPEC06 (based on SPEC06 c++ - mix of FP and Int tests)
  - Shown to scale well for LHC experiments
  - Simple conversion factor
  - Sites will benchmark existing capacity; vendors must run this benchmark suite (simple to run)
  - Process underway to convert requirements/pledges, and accounting
- **Automated gathering of installed capacity**
  - Process agreed between all parties – will be put in place to allow better understanding of available capacity; changes in information system will also improve normalisation between sites



# Validating the data...

## GridMap – Visualizing the "State" of the Grid

NorthernEurope	GermanySwitzerland	Italy	CERN
NO-NORGRID-T2	FZK-LCG2	UNI-DORTMUND	CERN-PROD
NDGF-T1	UNI-DORTMUND	Wuppertalprod	BEIJINGSFU
NIKHEF-ELPR	UNI-FREIBURG	MPPMU	CBVI/ALE
SE-SNIC-T2	DESY-HH	UNI-SIEGEN-HEP	SouthEasternEurope
UKI	France	SWITCH	AsiaPacific
UKI-SCOTGRID	IN2P3-CC-T2	IFCA-LCG2	Russia
UKI-LT2-IC	IN2P3-CC		
UKI-LT2-C			
RAL-LCG2			

**GermanySwitzerland**

**Total CPUs: 17928**

- FZK-LCG2: 5877
- UNI-DORTMUND: 2200
- DESY-HH: 2158
- RWTH-Aachen: 2024
- GoeGrid: 1680
- CSCS-LCG2: 960
- wuppertalprod: 928
- MPPMU: 800
- UNI-FREIBURG: 612
- DESY-ZN: 450
- SCAI: 108
- ITWM: 64
- MPI-K: 32
- UNI-BONN: 20
- GSI-LCG2: 9
- UNI-SIEGEN-HEP: 4
- SWITCH: 2
- UNI-KARLSRUHE: not in BDII
- LRZ-LMU: not in BDII

### Topology View

regions tiers pps all

sitenames  OSG sites

Size by:

CPU (GStat)

use historical CPU numbers

CPU (BDII) Running Jobs

use VOView information  size by SI2k

### SAM Results

Virtual Organization:

OPS Alice Atlas CMS

LHCb

Services:

Site CE SRMv2 sBDII

more..

Current Status:

latest SAM test results

Historical Availability:

hourly daily weekly monthly

**Size of site rectangles is number of CPUs from BDII.**  
**Certified Production sites, grouped by regions.**





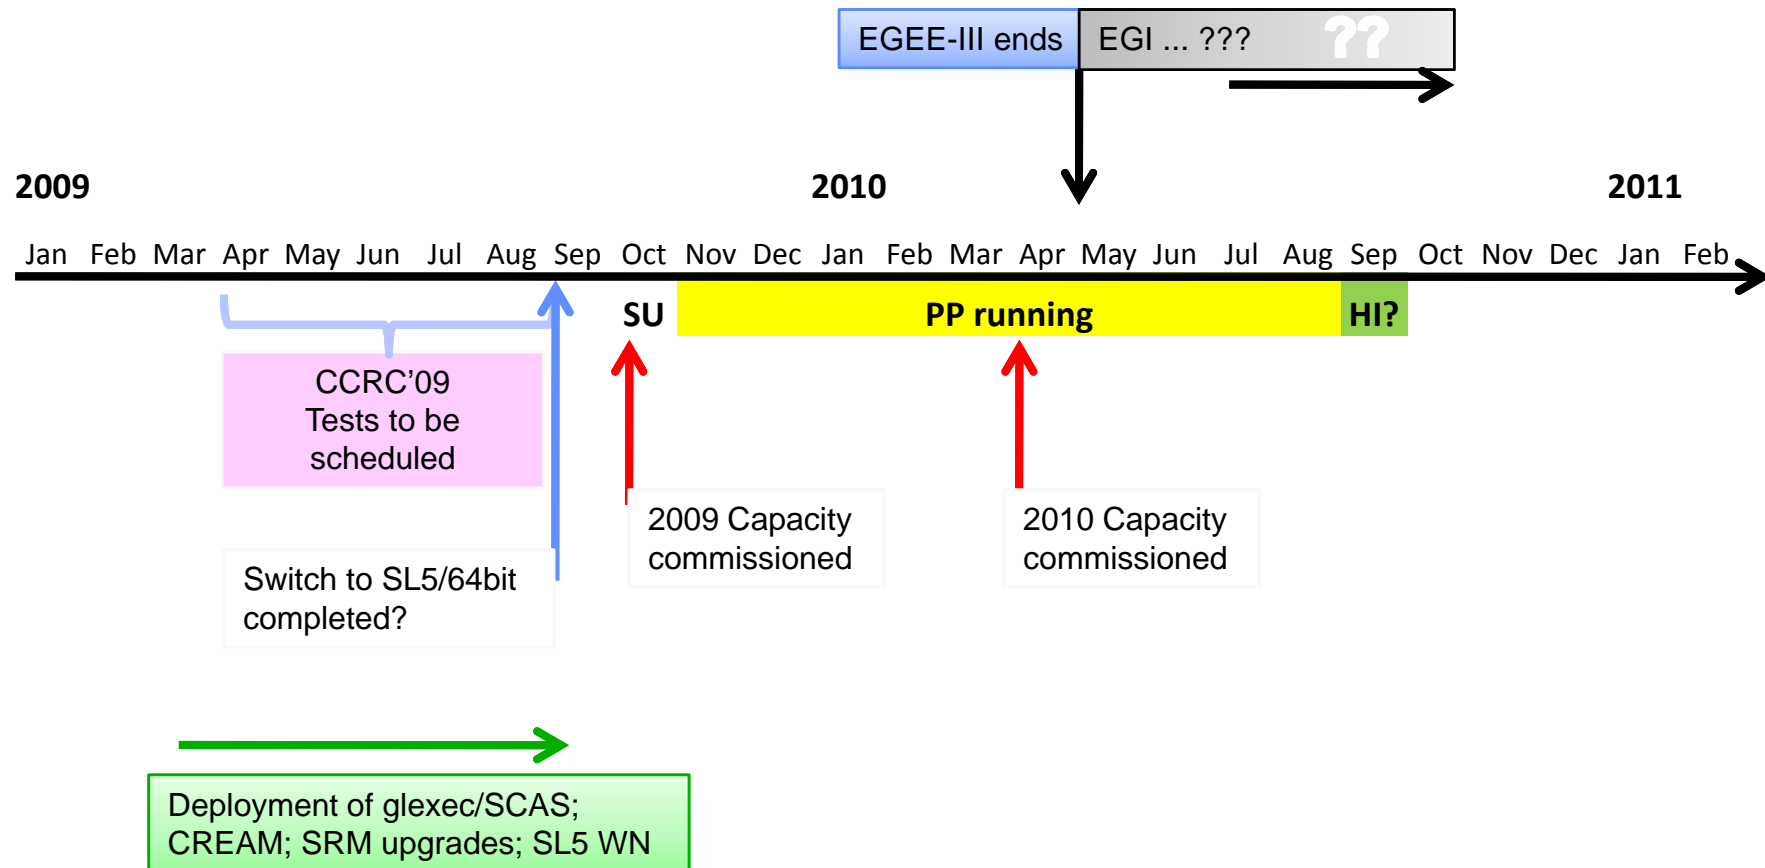
## CCRC'09 (revisited...)

- Originally a large combined test series did not fit with experiments' own testing schedules
- But,
- Tier 1s are concerned that we have not seen several/all experiments together testing tape recall/reprocessing at nominal rates, (and now we know we will need to do this while writing raw data to tape)
- All are concerned that we have not seen large scale tests of analysis
- LHCC mini-review conclusion:
  - "Recommend that there is a CCRC'09 in some form:"
    - At least CMS+ATLAS – but preferably with all 4 experiments
    - Testing reprocessing at Tier 1s (recall from tape) and massive/chaotic user analysis
    - Need metrics with which to evaluate this

**We urgently need to agree how this is done**



# WLCG timeline 2009-2010





# Planning for 2010 (end of EGEE)

- A final draft of the EGI blueprint has been produced (January)
  - Process will be discussed in the Overview Board next week
    - Document how the countries (Tier 1 + Tier 2) will provide the services and support needed for WLCG
      - Either as part of their NGI
      - Specific contribution
      - Must be no break in service at the end of EGEE-III
- EGEE-III have transition planning
- The Tier 0 is probably in a reasonable position – current planning does not rely on external funding; but the capability will be strictly limited to core WLCG Tier 0/CAF tasks
  - The location of the EGI.org is being studied now – bids have been received
    - Decision 1<sup>st</sup> week of March at EGEE User Forum ?
  - Still not clear how many NGIs will really exist in 2010 to support this



# Pending issues for 2009

- Plan to have visits of Tier 1 sites – to understand service issues
  - MSS
  - Databases – seems to be often a source of problems
  - Share and spread knowledge of running reliable services
- SRM performance
  - Need good testing of Tier 1 tape recall/reprocessing, together with Tier 1 tape writing – for several experiments together
    - Encapsulated tests?
  - Clear need for storage “reassessment” (tbd...) from yesterday’s discussion
- Data access with many users for analysis – need to see experiment testing of analysis
- Transition plan for 2010 – to cover services today provided by EGEE
  - May be short or long term – but is probably going to be needed



# Summary

- CCRC'08 was successful
  - Almost all experiments' and service targets were achieved
    - Exception: user analysis with 100's of users; T1 reprocessing at full scale
- Service has continued to be used at significant levels
  - Main focus is on improving service reliability – especially storage systems
- Important that resource ramp-up for 2009/10 continues:
  - Delay allows for more effective purchases in some cases
  - Resource procurements/installations were significantly delayed in 2008
  - Must be ready for the accelerator start-up, even if resources are today not saturated
- Planning for future – Tier 0/CAF and European Grid infrastructure – ongoing



# Other topics

- Building the WLCG community –
  - Common projects? Esp. of interest to Tier 2s
    - DPM, Nagios+tests/probes
    - Future support for sites e.g. In Latin America when EGEE/other EU projects end?
    - Training site managers on middleware, services, tools, best practices, etc(workshops, ???)
    - ...
- Communications
  - How do Tier 2s get/stay informed?
    - Suggestions?