WLCG Status Report

CERN-RRB-2009-040

28th April, 2009

Computing Resource Review Board

Ian Bird
LCG Project Leader

Agenda

- WLCG service status
- Roadmap – STEP'09
- Planning for 2010 (EGI etc)
- Re-assessment of experiment requirements

WLCG Service

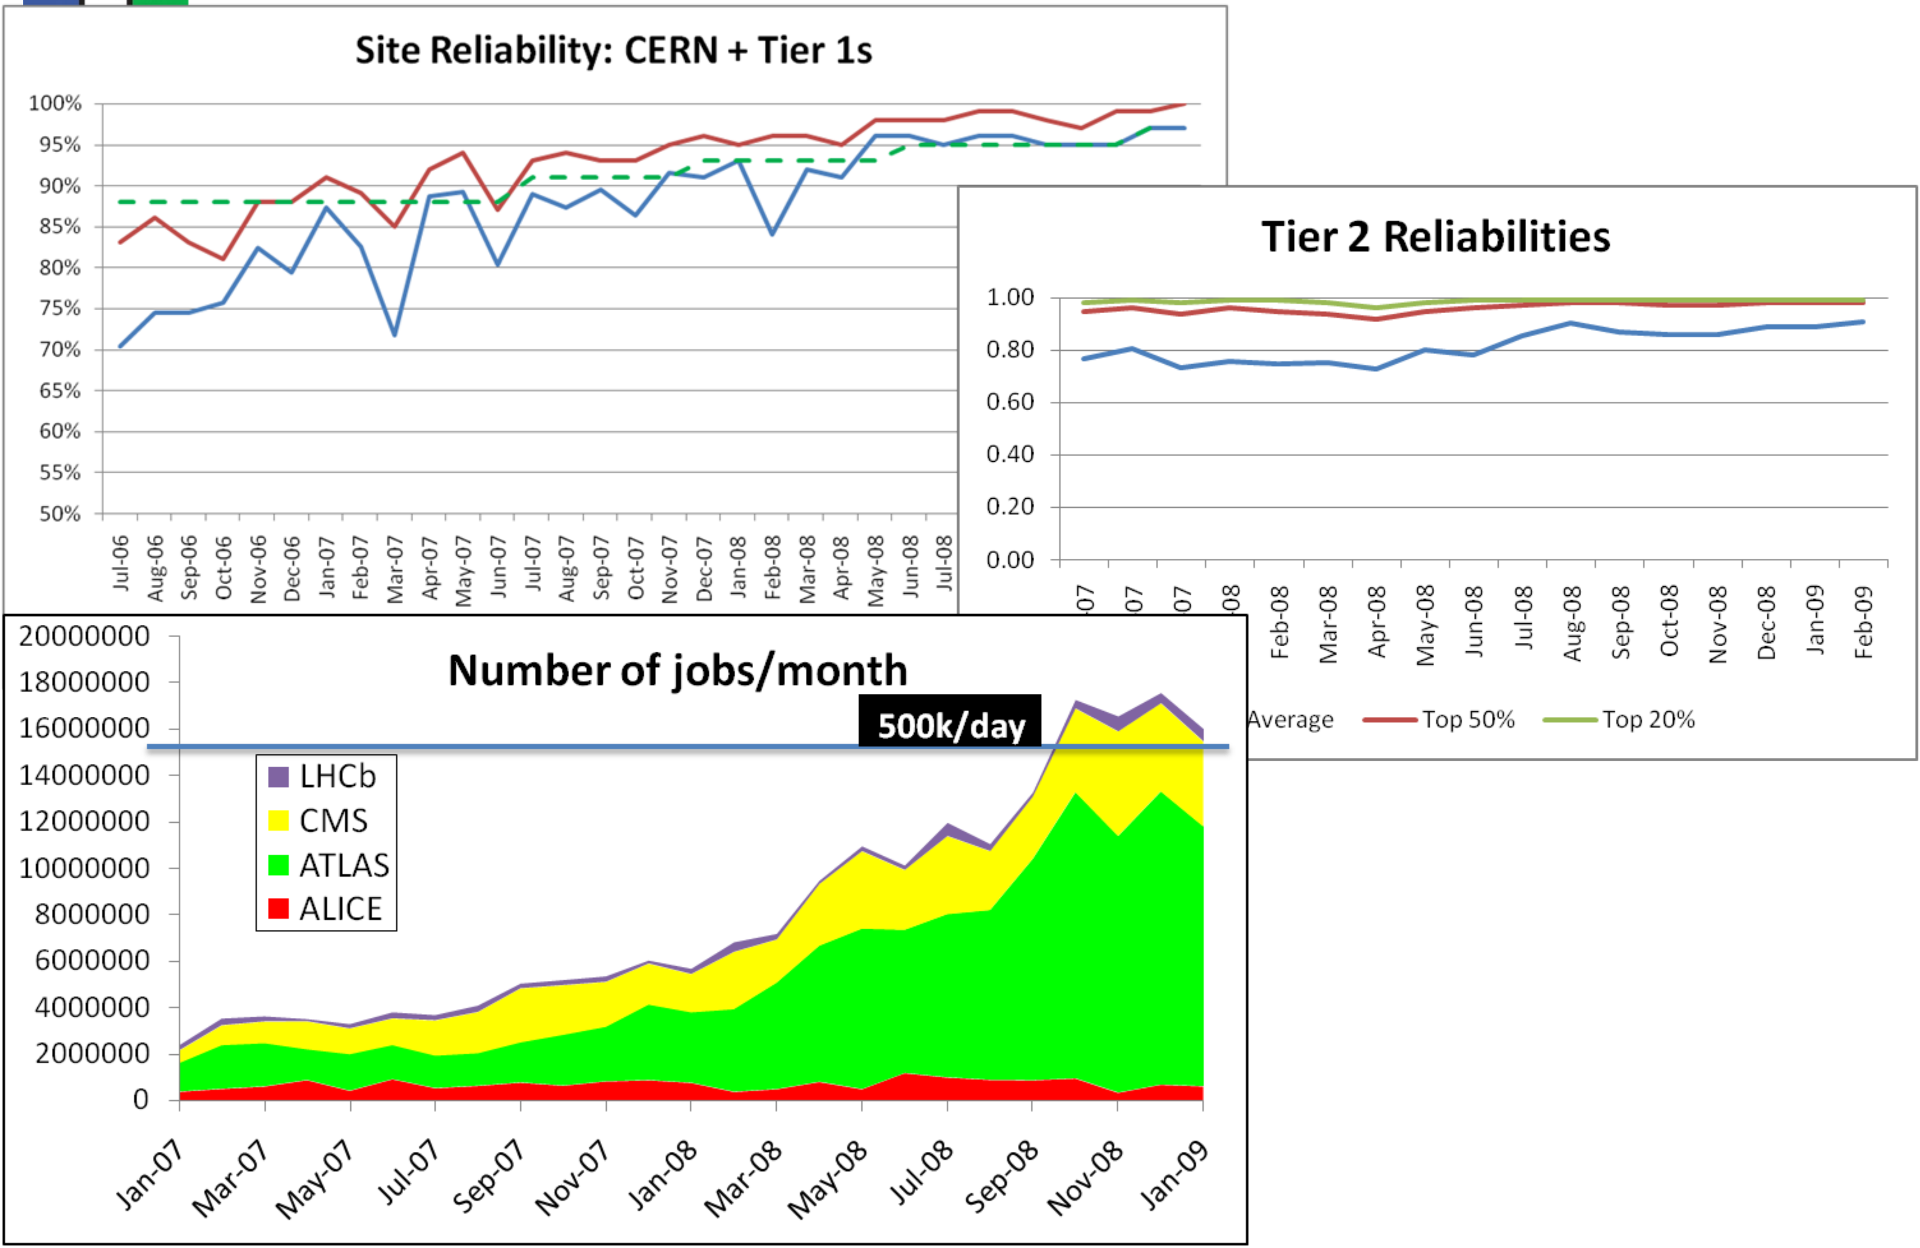

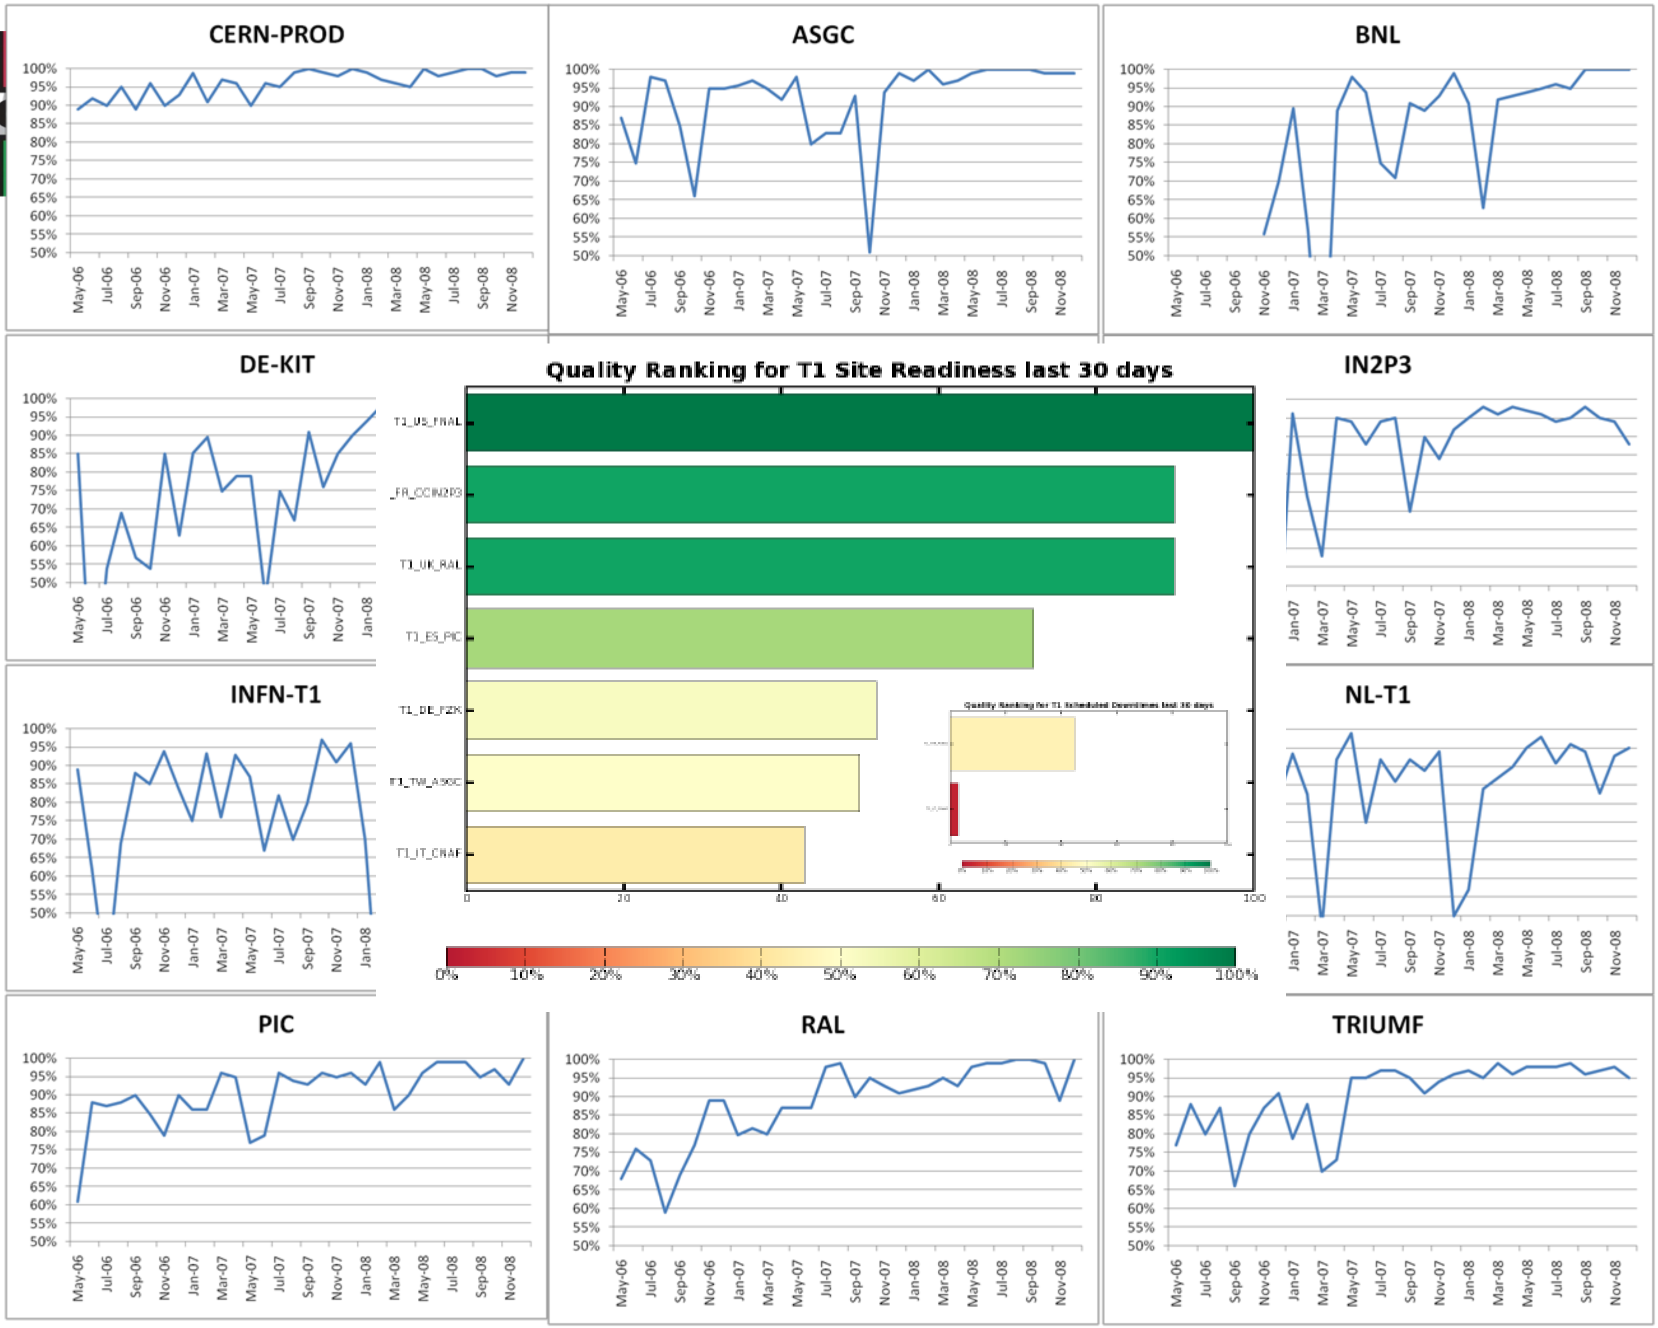
CERN-PROD
ASGC
BNL
DE-KIT
Quality Ranking for T1 Site Readiness last 30 days
T1_US_FNAL
_FR_CCIN2P3
T1_UK_RAL
T1_ES_PIC
T1_DE_FZK
T1_TW_ASGC
T1_IT_CNAF
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%
IN2P3
INFN-T1
NL-T1
PIC
RAL
TRIUMF

Federation	Reli-ability	Avail-ability
T2_US_Caltech	100 %	100 %
T2_US_UCSD	100 %	100 %
US-NET2	99 %	99 %
T2_US_Wisconsin	99 %	95 %
JP-Tokyo-ATLAS-T2	99 %	99 %
FR-IN2P3-SUBATECH	99 %	99 %
ES-ATLAS-T2	99 %	95 %
FR-IN2P3-LAPP	99 %	98 %
FR-IN2P3-LPC	99 %	99 %
T2_US_MIT	98 %	98 %
T2_US_Purdue	98 %	93 %
CH-CHIPP-CSCS	98 %	97 %
US-SWT2	98 %	98 %
HU-HGCC-T2	98 %	98 %
AT-HEPHY-VIENNA-UIBK	97 %	87 %
BE-TIER2	97 %	91 %
US-MWT2	97 %	97 %
UK-NorthGrid	97 %	97 %
DE-DESY-ATLAS-T2	97 %	97 %
US-WT2	97 %	97 %
CN-IHEP	97 %	96 %
T2_US_Florida	97 %	97 %
CZ-Prague-T2	97 %	97 %
PT-LIP-LCG-Tier2	97 %	93 %
FR-IN2P3-CC-T2	97 %	89 %
UK-ScotGrid	96 %	96 %
ES-LHCb-T2	96 %	96 %
UK-London-Tier2	96 %	93 %
KR-KISTI-T2	95 %	95 %
RO-LCG	95 %	95 %
PL-TIER2-WLCG	95 %	95 %

Federation	Reli-ability	Avail-ability
UK-SouthGrid	95 %	94 %
T2_US_Nebraska	95 %	95 %
FR-GRIF	95 %	92 %
ES-CMS-T2	94 %	89 %
FR-IN2P3-IPHC	93 %	93 %
DE-DESY-RWTH-CMS-T2	93 %	91 %
US-AGLT2	93 %	93 %
IN-DAE-KOLKATA-TIER2	93 %	84 %
RU-RDIG	93 %	92 %
IT-ALICE-federation	92 %	81 %
IT-ATLAS-federation	92 %	81 %
IT-CMS-federation	92 %	81 %
IT-LHCb-federation	92 %	81 %
DE-MCAT	90 %	84 %
FI-HIP-T2	87 %	80 %
AU-ATLAS	87 %	50 %
DE-FREIBURGWUPPERTAL	86 %	78 %
SE-SNIC-T2	86 %	85 %
NO-NORDGRID-T2	85 %	84 %
IN-INDIACMS-TIFR	85 %	83 %
CA-WEST-T2	84 %	83 %
EE-NICPB	84 %	81 %
TR-Tier2-federation	83 %	83 %
DE-GSI	79 %	28 %
KR-KNU-T2	76 %	73 %
IL-HEPTier-2	75 %	75 %
CA-EAST-T2	34 %	15 %
PK-CMS-T2	0 %	0 %
SI-SiGNET	0 %	0 %
TW-FTT-T2	0 %	0 %
UA-Tier2-Federation	N/A	N/A

- Tier 2 reliabilities
- Goal to have all >95%
- Significant improvements
- Only 1 non-reporting federation

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 time_site_is_available / (total_time - time_site_is_scheduled_down)
Target reliability for each site is 95 % and Target for 8 best sites is 97 % from June, 2008

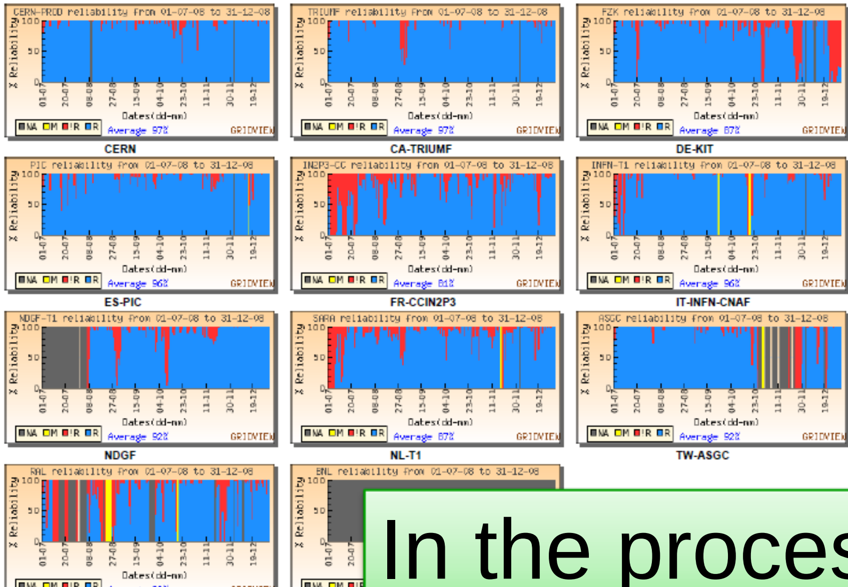

Reliability of WLCG Tier-1 Sites + CERN for CMS July 2008 - December 2008

Data from SAM Monitoring. Plots show Reliability for last 6 Months
Reliability is calculated as time_site_is_available / (total_time - time_site_is_scheduled_down)
Target reliability for each site is 95 % and Target for 8 best sites is 97 % from June, 2008

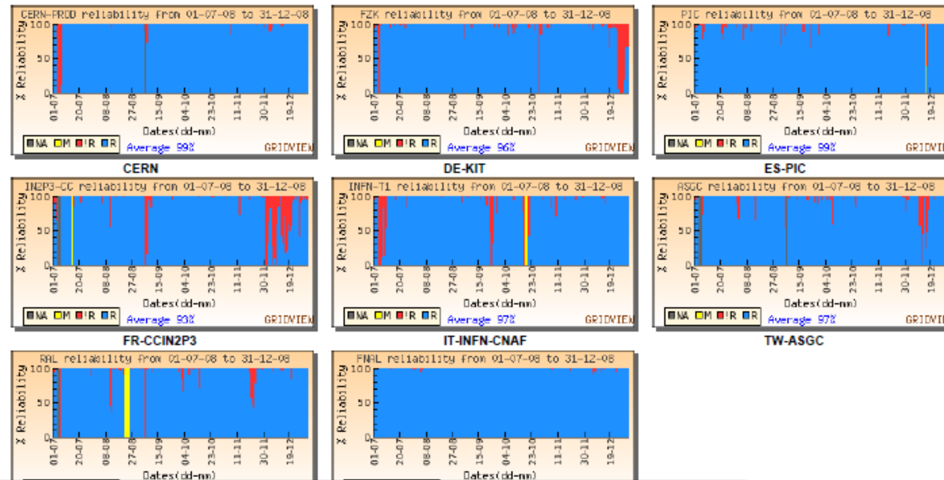

In the process of being validated

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 time_site_is_available / (total_time - time_site_is_scheduled_down)
Target reliability for each site is 95 % and Target for 8 best sites is 97 % from June, 2008

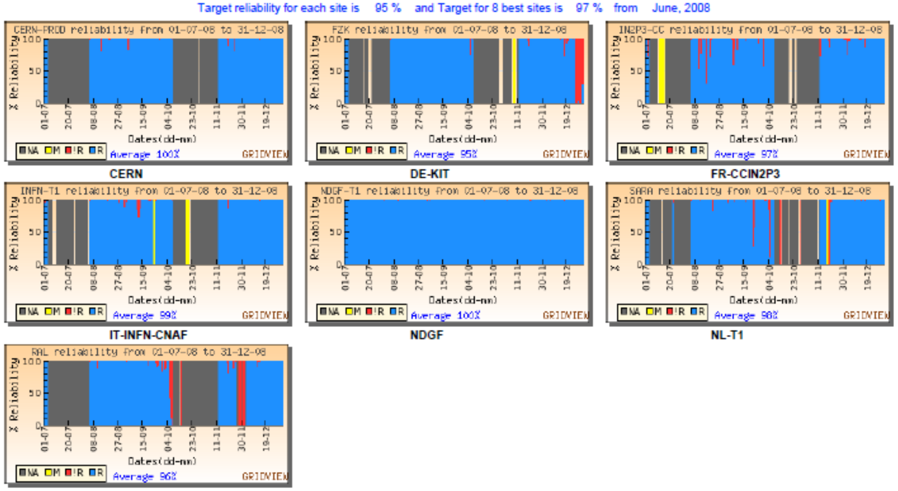

- December 2008

Data from SAM Monitoring. Plots show Reliability for last 6 Months
Reliability is calculated as time_site_is_available / (total_time - time_site_is_scheduled_down)
Target reliability for each site is 95 % and Target for 8 best sites is 97 % from June, 2008

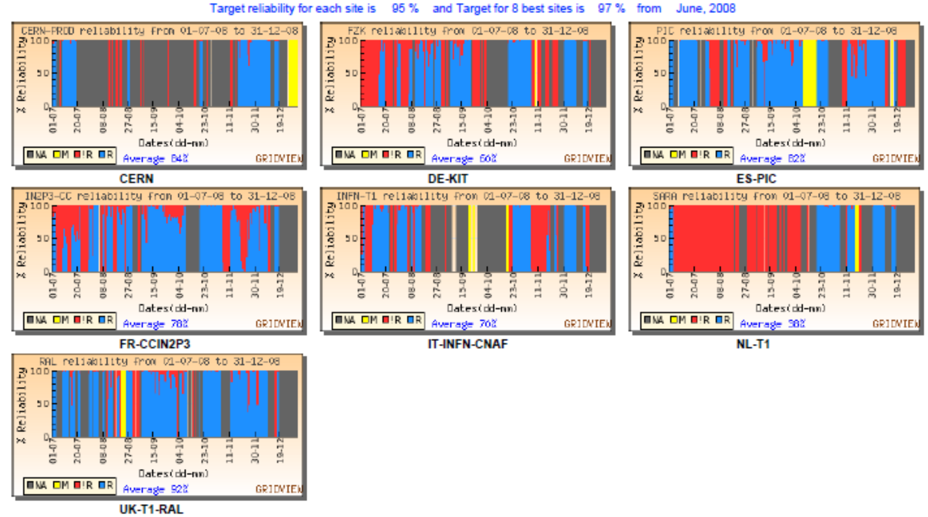

Total

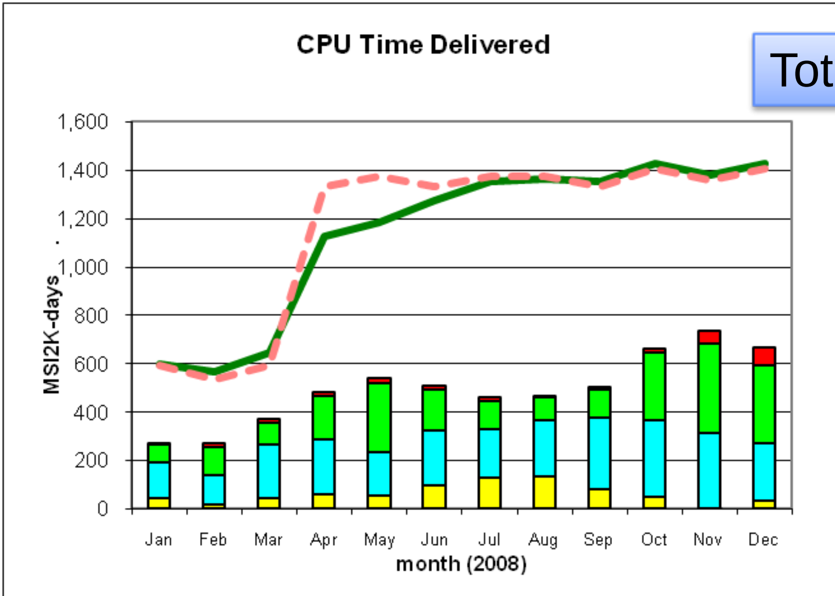

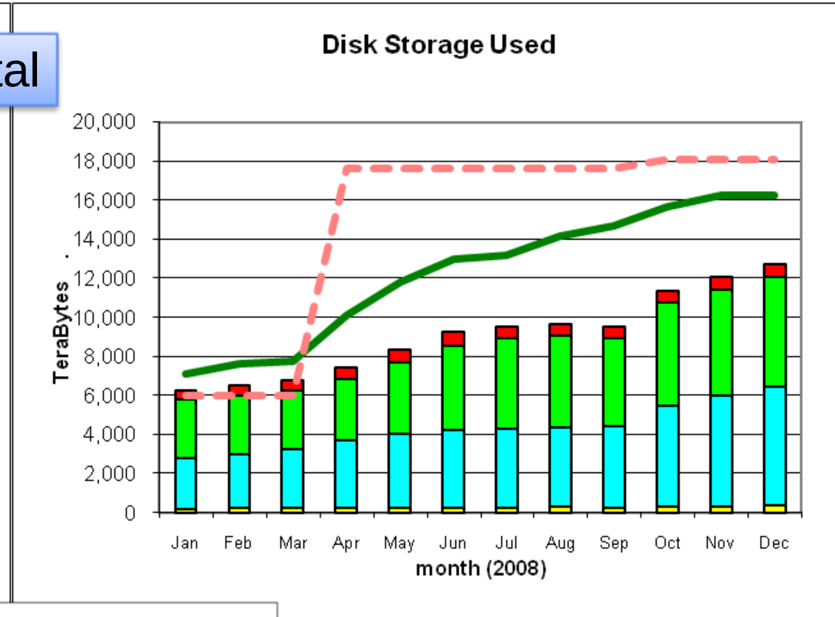

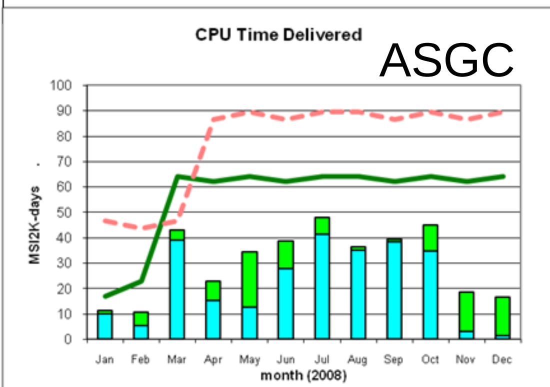

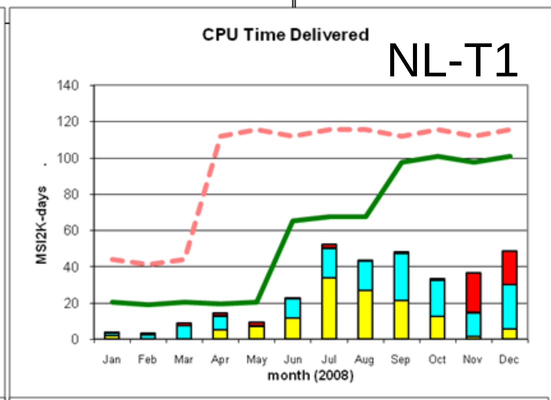

- Many problems to ramp up resources
  - Delays in procurements
  - Faulty equipment
  - Lack of power & planning

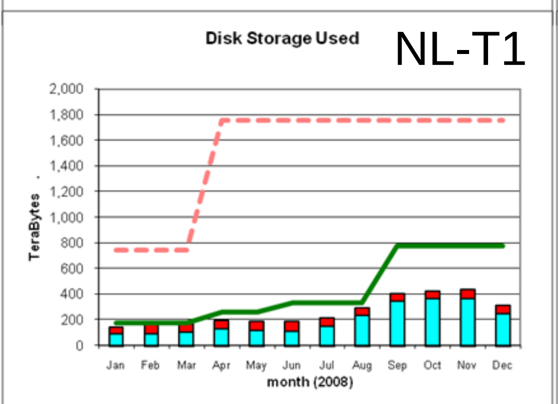

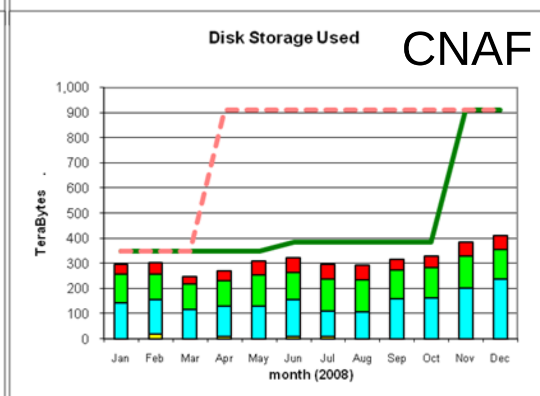

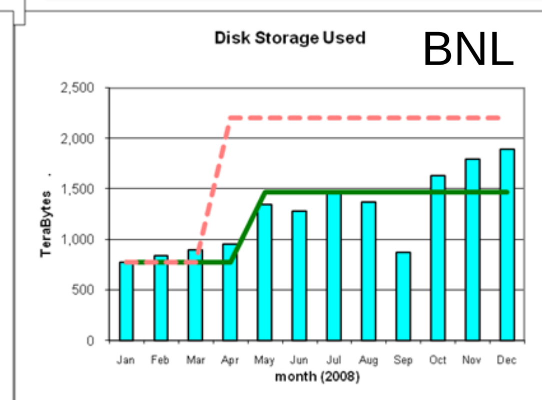

Serious Incidents.

In last six months

- Castor – ASGC, CERN, CNAF, RAL
- dCache – FZK, IN2P3, NL-T1
- Oracle – ASGC, RAL
- Power – ASGC, PIC, NL-T1, CNAF
- Cooling- CERN, IN2P3
- Network- CNAF, PIC, BNL, CERN
- Other – CNAF, RAL, NL-T1,
- Fire – ASGC

Tier1s **will** be down. Experiment models should cope.

Improving reliability

- Simple actions
  - Ensure sites have sufficient local monitoring; including now the grid service tests/results from SAM and experiments
  - Ensure the response to alarms/tickets works and is appropriate – test it
  - Follow up on SIRs – does your site potentially have the same problem???
  - If you have a problem be honest about what went wrong – so everyone can learn
- Workshops
  - To share experience and knowledge on how to run reliable/fault tolerant services
    - WLCG, HEPiX, etc.
- Visits
  - Suggested that a team visits all Tier 1s (again!) to try and spread expertise ...

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
  - Resource requests now given in kHS06
- 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...

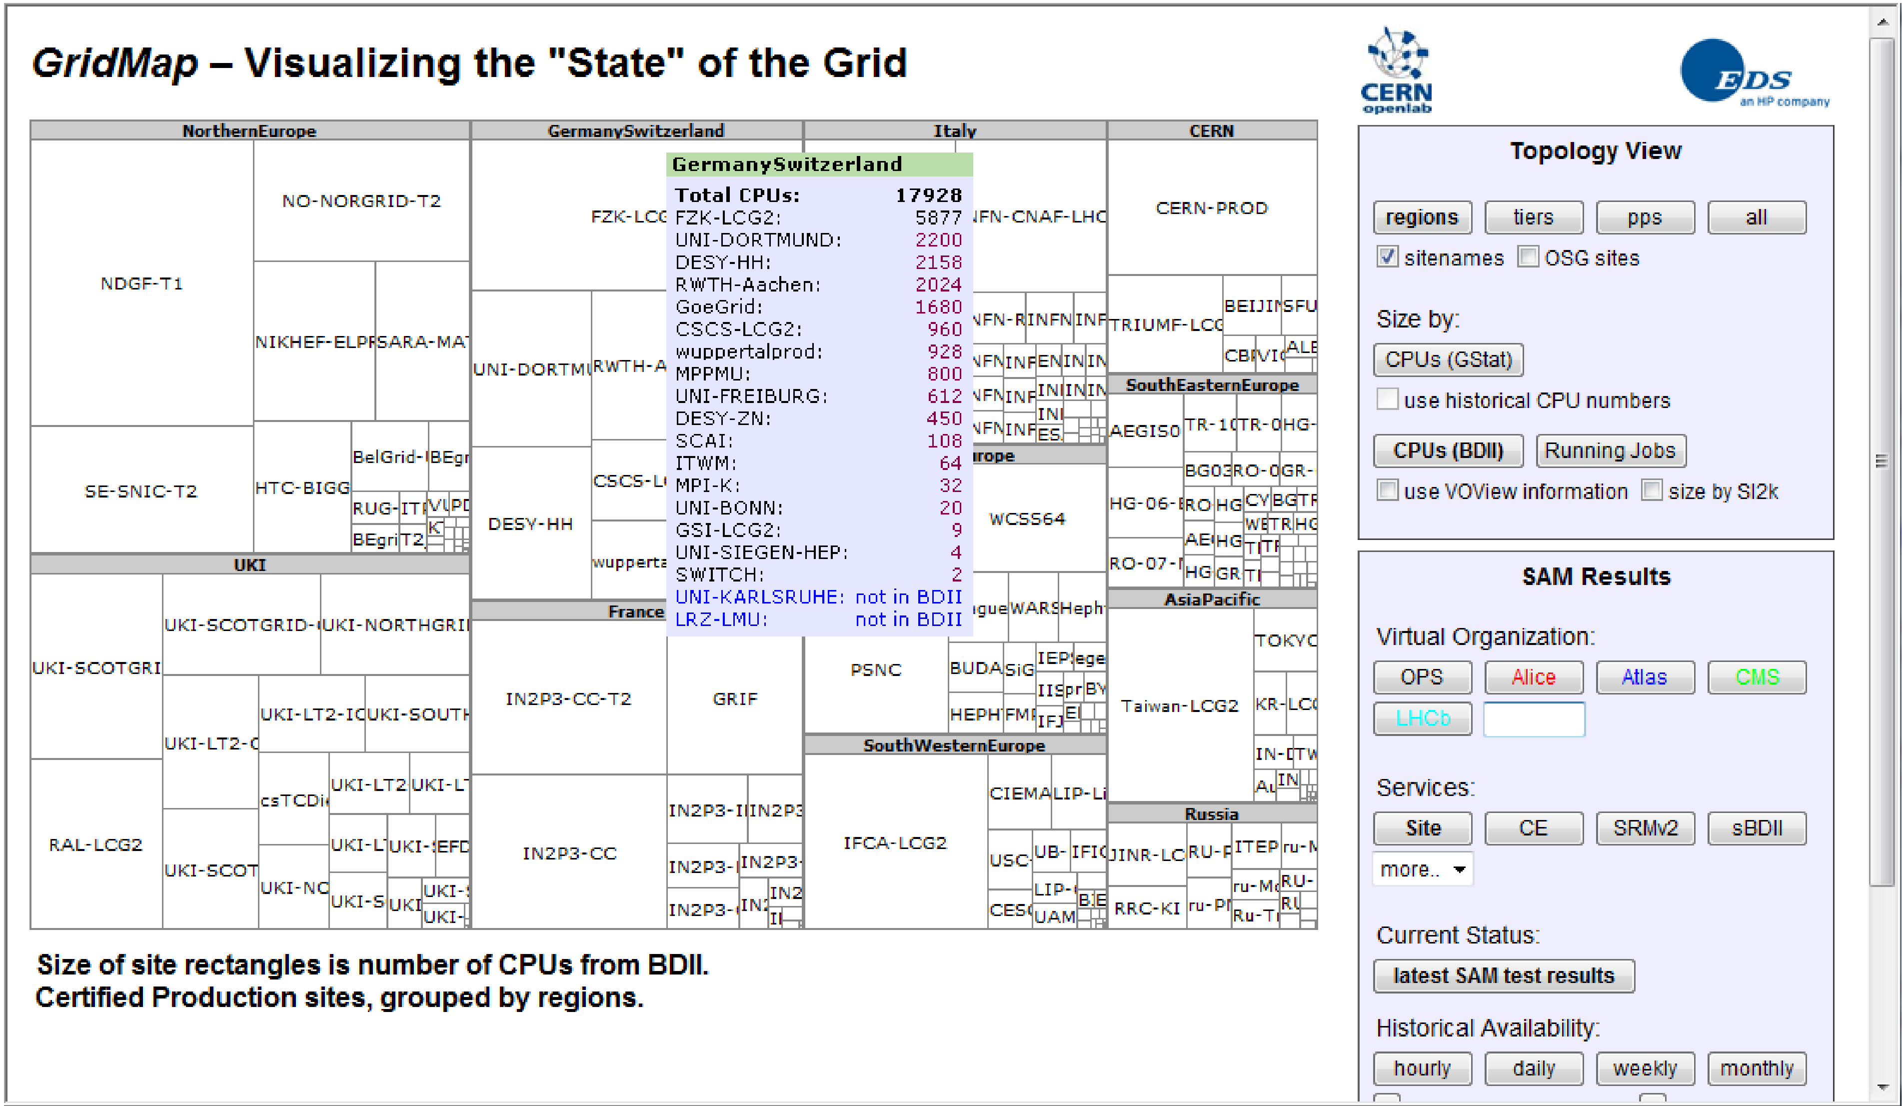

Size of site rectangles is number of CPUs from BDII.
Certified Production sites, grouped by regions.

WLCG timeline 2009-2010

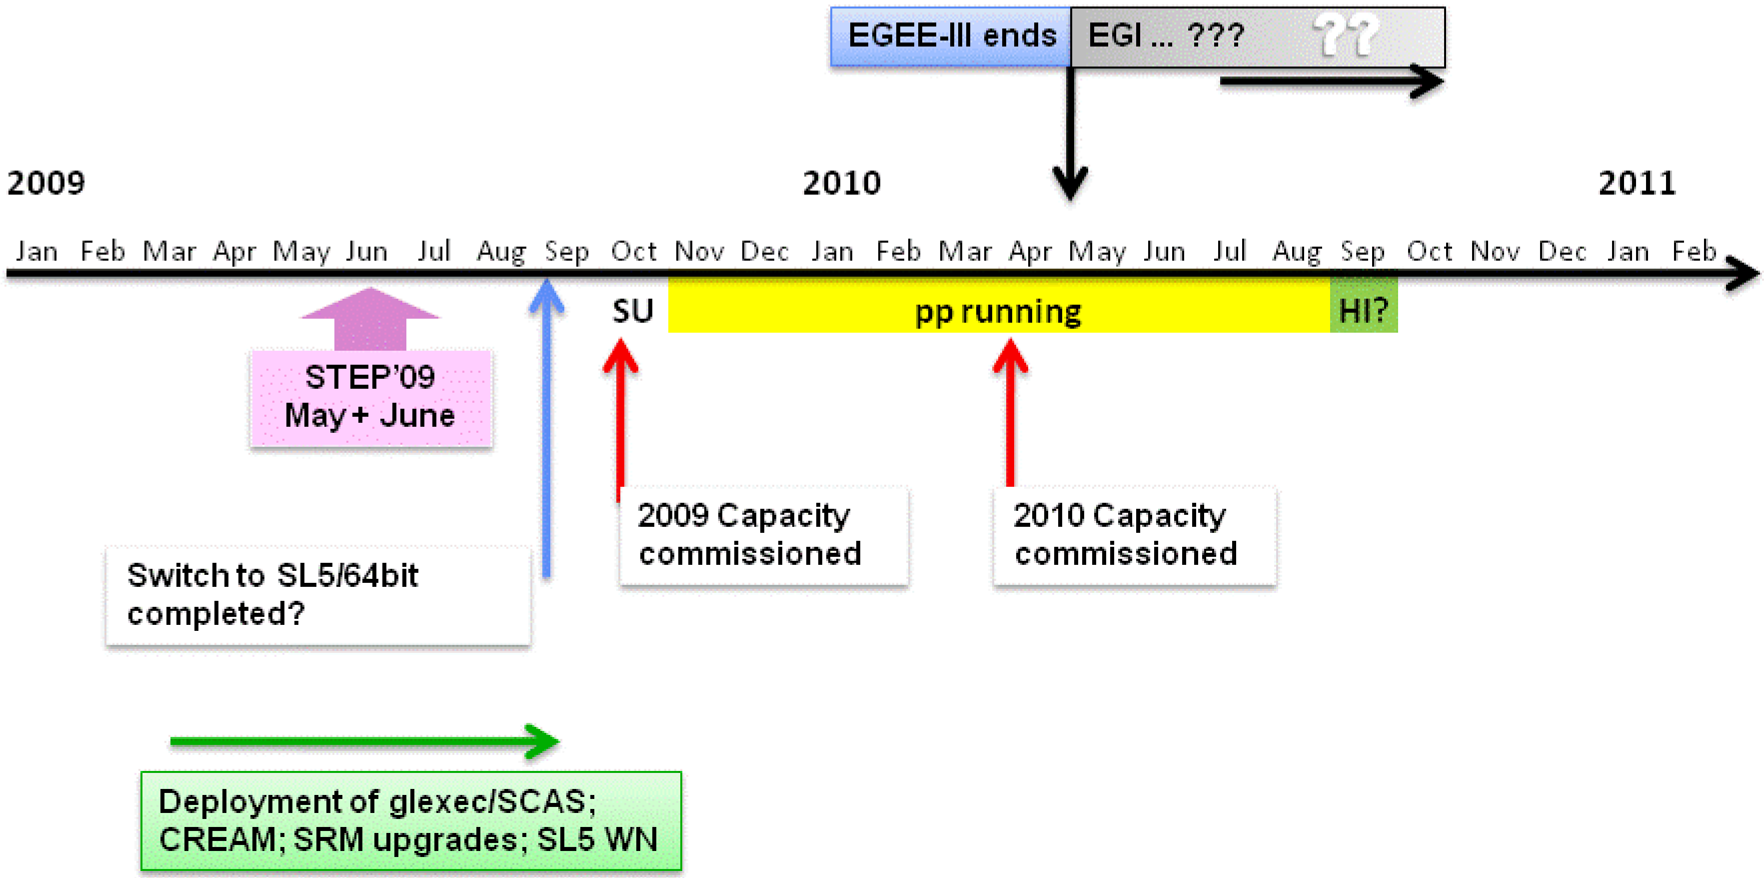

Ian.Bird@cern.ch

STEP'09

- While CCRC'08 was seen as a success, it did not fully test all of the essential parts of the computing models
  - Tape recall for reprocessing at Tier 1s – with >1 experiment at full rates
  - Large scale analysis
- Recommended by LHCC that there should be such tests, but difficulty in scheduling between experiments
  - Their schedules driven by detector related work
- At WLCG workshop in Prague – agreed that the testing schedules could be aligned
- "Scale Testing for the Experimental Programme in 2009" – STEP'09
  - (i.e. Will probably have increasing scale tests in future years)
  - May – June 2009
  - Experiments have full programmes of testing, but will co-schedule tape-recall/reprocessing at Tier 1s, and analysis scenarios
  - Key metrics being agreed now

Planning for after EGEE: EGI

- Status
  - Final blueprint published end December
  - EGEE transition plan produced, helped in final blueprint
    - While cannot fully implement transition, clarifies expected state at end of project
  - EGI.org+NGIs will take over the infrastructure – transition plan

  - EGI_DS Policy Board has selected Amsterdam as location for EGI.org
  - Establishment of organisation:
    - Council of NGIs – with an initial MoU (LoI in first instance)
    - EGI.org must appoint Director and for teams to develop transition plans etc.
  - WLCG has made statement of support for the process and willingness to work with EGI.org and NGIs; expects to participate via the User Forum

EGI - timescales

- Timescales as presently understood
  - Early April: MoU and LoI available
  - End April: LoI signed by interested NGIs
  - May 6: (proto)EGI Council established
    - NGIs signing LoI are constituents
    - The EGI Project(s) team confirmed; this includes the project director(s) identification/confirmation (the person who will lead the team(s))
  - June: MoU signed, (full)EGI Council setup
    - “The Transition towards EGI” Deliverable published
  - March—May: EGI.org setup preparation
    - Includes search for EGI.org director and identification of EGI.org key personnel
  - June: EGI.org director appointed/identified
  - September: EGI.org setup at the latest (one month before the call closure)
  - July—December: EGI Project(s) preparation and submission
  - The MoU signing will continue after June, but the “latecomers” may not have direct influence on the composition of project preparation team nor on the selection of EGI.org director

Updated resource requirements

- Based on the presently understood LHC schedule

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 '	SH	SH	SH	SH	SH	SH	SH	SH	SU	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	PH	SH	SH	SH	SH	SH	SH
	44 weeks physics possible																									

- For planning purposes we assume
  - 2 resource periods (although no break between them)
    - “2009” Oct’09 → March’10
    - “2010” April’10 → March’11 (as before)
  - For data taking:
    - Apr’09 – Sep’09: no LHC (simulation and cosmics)
    - Oct’09 – Mar’10: 1.7x10^6 sec of physics
    - Apr’10 – Oct’10: 4.3x10^6 sec of physics
    - Nov’10 – Mar’11: LHC shutdown (simulation, reprocessing, etc)
  - Energy is limited to 5+5 TeV
  - There will be a heavy-ion run at the end of 2010

General comments

- Overall there is less LHC live time in this period than was anticipated for 2009+2010, but #events is only ~x2 less
- However,
  - We must ensure that the computing is not a limiting factor when data comes
    - See LHCC conclusions of WLCG mini-review in February
  - Significant effort is going into detector understanding now using cosmic ray data
- Early in 2009 we relaxed the requirement to have the 2009 resources in place by April
  - Although many of the (Tier 1) resources are actually in place now
  - In some cases this allows delayed procurement for better equipment
- Will now need to install new resources while data taking
- Intend to eventually provide a profile of ramp-up of resources (quarterly?) – but for this discussion present only the total needs for the 2 resource periods
  - Helps with installation schedules

Comparisons

- For each experiment (in the following tables):
  - Updated requirement for 2009 and 2010 compared with existing 2009 pledge and old 2010 requirement (since we do not have the split between experiments for pledges after 2009)
- Overall requirements
  - Compare 2009, 2010 new requirements with existing pledges
- The new requirements have not been reviewed by the LHCC, the C-RSG has had only a few days to look at them
- The pledges do not take into account:
  - Change in INFN planning, nor delay at NL-T1, and others where 2008 pledges not yet fully installed

ATLAS

ATLAS	2009 req	2009 pledge	2010 req	Old 2010 req
CERN CPU	57	26.5	67	68
CERN disk	3.7	2.075	5.1	5.25
CERN tape	7.8	6.21	9.9	14.6
T1 CPU	90	120.9	227	234
T1 disk	24	19.86	36.7	41.3
T1 tape	11.3	14.72	14.8	22.7
T2 CPU	108	114	240	242
T2 disk	13.3	11.2	24.8	24.8

Requirement >10% more/less than **pledge/requirement**

- Cosmic ray data in Q309 will produce 1.2PB (same as Aug-Nov 08)
- In 6x10^6 sec will collect 1.2x10^9 events ➔ 2PB raw
- Raw stored on disk at T1s for a few weeks
- Plan for 990M full sim events and 2200M fast sim events
- CERN request was updated last Aug and was seen by RSG

- Generally new requirements <= old requirements (except at CERN)
- Provide resource needs profile by quarter (see document)
- NB. The August 2008 request for 2009 while agreed by the RSG has never been validated by LHCC

CMS

CMS	2009 req	2009 pledge	2010 req	Old 2010 req
CERN CPU	48.1	54.8	112.9	115.2
CERN disk	1.9	2.5	4.6	3.8
CERN tape	9.5	9.3	15.3	14.3
T1 CPU	53.5	63.7	119	139
T1 disk	6.5	8.4	14.1	15.4
T1 tape	10.5	16	21.6	23.2
T2 CPU	54.1	116	209.6	306
T2 disk	5	8.4	11.3	7.6

Requirement >10% more/less than **pledge/requirement**

- Model foresees 300Hz data taking rate ...
- ... and CPU times assume higher lumi in '10
  - recCPU: 100→200 HSO6.s
  - simCPU: 360→540 HSO6.s
- Changes
  - 3 re-reconstr in each '09, '10
  - 40% overlap in PD datasets
  - Added storage needs for '09 cosmics

- Tier 0:
  - Added 1 re-reco in each year
  - Capacity for express stream
  - Reco to finish in 2x runtime in '09
  - Monitoring + commissioning is now 25% of total (was 10%)

- Tier 1:
  - Finish '09 re-reco in 1 month (was spread over full year)
- Tier 2:
  - Require 1.5 more MC events than raw: sw changes and bug fixes
  - MC events produced in 8 months (can only start after Aug'09)

ALICE

ALICE	2009 req	2009 pledge	2010 req	Old 2010 req
CERN CPU	42.8	46.4	46.8	49.4
CERN disk	2.4	4.5	4.5	4.7
CERN tape	3.7	7.3	6.7	11.6
T1 CPU	42.8	40.9	102.4	94
T1 disk	4.3	3.9	9.9	12
T1 tape	5.9	6.2	11.6	19.7
T2 CPU	36	39.9	80.8	100
T2 disk	4.4	2.82	12.4	4.3

Requirement >10% more/less than **pledge/requirement**

- Will collect p-p data at ~maximum rate: 1.5x10^9 events at 300 Hz
  - Initial running will give luminosity required without special machine tuning – cleaner data for many physics topics
  - First pp run energy is important in interpolating results to full Pb-Pb energy
- Thus plan to collect large statistics pp in 2009-10
- Assume 1 month Pb-Pb at end of 2010

- Requests are within (or close to) existing '09 pledges except for Tier 2 disk
- For 2010 – don't know actual pledge for ALICE, but generally pledges are significantly lower than requirement. (so final column should be mostly pink for T1+T2!)

LHCb

LHCb	2009 req	2009 pledge	2010 req	Old 2010 req
CERN CPU	17	4.2	28	6.12
CERN disk	0.78	0.99	1.47	1.28
CERN tape	1.2	2.27	2.3	4.2
T1 CPU	31	20.2	49	27.36
T1 disk	2.8	2.7	4.4	3.25
T1 tape	1.3	3.2	2.9	5.86
T2 CPU	30	35.4	40	45.5
T2 disk	0.02	0.37	0.02	0.02

Requirement >10% more/less than **pledge/requirement**

- Uncertainty in running mode (pile up) ➔ add contingency on event sizes and simulation time
- 2009 Simulation with assumed running conditions
- Early data with loose trigger cuts and many reprocessing passes – alignment/calib+early physics
- 2010 – several reprocessing passes and many stripping passes
- Simulation over full period

- CERN increase due to need for fast feedback to detector of alignment/calibration + anticipation of local analysis use
- T1 CPU increase in 2010 due to more reprocessing
- T2 requirements decrease as less overall simulation needed

NB. Previously LHCb had presented integrated CPU needs – now here are shown the total capacity needed in each period – as for the other experiments

Summary

Summary	2009 req	2009 pledge	2010 req	Old 2010 req	2010 pledge
CERN CPU	164.9	131.9	254.7	238.7	213.6
CERN disk	8.78	10.07	15.67	15.03	13.4
CERN tape	22.2	25.1	34.2	44.7	43.1
T1 CPU	217.3	245.7	497.4	494.36	406.1
T1 disk	37.6	34.9	65.1	72	60.3
T1 tape	29	40.12	50.9	71.46	65.9
T2 CPU	228.1	305.3	570.4	693.5	475.8
T2 disk	22.72	22.79	48.52	36.72	35.2

Requirement >10% more/less than **pledge**

Live time is 1/3 of that anticipated for 2009/10 – so why so many resources needed?

- This is the first year of data taking and is a critical opportunity
  - Must be able to analyse and react quickly – provide feedback to the detectors during data taking
- Machine profile is different from the original plans – long run followed by a long shutdown
  - Must ensure that problems are resolved as soon as possible
  - There is now competition – fast reaction is essential
- We have experience now in executing almost all of the computing models; and with real data (from cosmics)
  - Important lessons have been learned and the models refined
  - Analysis from disk is essential – the idea of allowing tape access for many users is not supportable
  - Last year has shown more Tier 0/CAF resources are needed
- Use of cosmics (and real data) in 2008/9 – experiments have made huge progress in understanding the detectors – would otherwise have been done with beam
  - Can thus more rapidly focus on extracting physics when collisions arrive

Implications:

- CPU needs driven by the instantaneous requirement, e.g. ability to react rapidly and provide early feedback during data taking
  - Experiments will take data at the intended rate, ~independent of luminosity, although complexity (and CPU) increases with luminosity
  - Need to provide rapid analysis and feedback; anticipate additional reconstruction passes
- Storage needs (disk) partly correlated with CPU needs, partly with longer term:
  - #events (#files) driven by trigger rate – thus drives storage needs
  - Need sufficient disk to allow the rapid feedback analysis
  - Access to tape for analysis is not feasible – must have sufficient disk to enable all of the urgent analyses that must be done – and support the number of people involved
  - In this first year data must be kept on disk for some time for (re-)analysis
- Cosmic data has been used to understand the detectors in preparation for real data
  - This data is invaluable in understanding the history and trends of the detectors behaviour and cannot be simply discarded
  - Amount is significant (e.g. ATLAS ~2 PB)

Implications cont...

- Most of the experiments actually provide a resource profile over the full period April 2009 – March 2010; by quarters
  - This can help in scheduling installations – and even purchases in some cases
    - But experience in procuring/testing/installing in time is not good
  - Much of the CPU is needed earlier, but storage capacity can be ramped to a certain extent
  - Most of the high quality useful data will come in the last quarter of running – everything must be done to prepare for that

Summary

- WLCG service supports ever increasing workloads
- Significant incidents continue at fairly high level
  - Not all can be avoided
- Must take lessons from the reports and adapt
- Service reviews and STEP'09 will occupy time until LHC startup
- Requirements for 2009/10 have been reassessed in view of:
  - New LHC schedule and anticipated running profile
  - Experience with the computing models in full scale use in the last year and with cosmic data for several months
- Overall increased needs for CPU and disk (esp in 2010); reduced need for tape
- But: this is the first year of data taking which will be followed by a long shutdown – only opportunity to get it right
  - Must ensure that the computing resources are not a limiting factor compared with the huge investment in the accelerator and detectors