WLCG Status Report

31st March, 2008
Project Overview Board

Ian Bird
LCG Project Leader

MoU Status

- Since the December meeting of the OB:
  - T1: Norway and Sweden have signed
  - T2: Hungary (ALICE, CMS), Germany (ATLAS T2), Norway (ATLAS), Sweden (ALICE, ATLAS), Turkey (ATLAS, CMS)

- Still open:
  - 2 new Canadian Tier 2 federations will sign soon
  - Brazil ready to sign if delegation of Funding Agency signature is acceptable to OB
  - Timetables for Austria (ATLAS, CMS), and Czech Republic (15/4/08)?
  - Formal US contribution to ALICE?

- Signatures (32 countries)
  - Australia, Belgium, Canada, China, Denmark, Estonia, Finland, France, Germany, Hungary, India, Israel, Italy, Japan, JINR (Dubna), Korea, Netherlands, Norway, Pakistan, Poland, Portugal, Romania, Russia, Slovenia, Spain, Sweden, Switzerland, Taipei, Turkey, UK, Ukraine, USA.

Level 1 Milestones

27 Mar 08

Done	Late <1 mon	Late >1 mon

colour coding

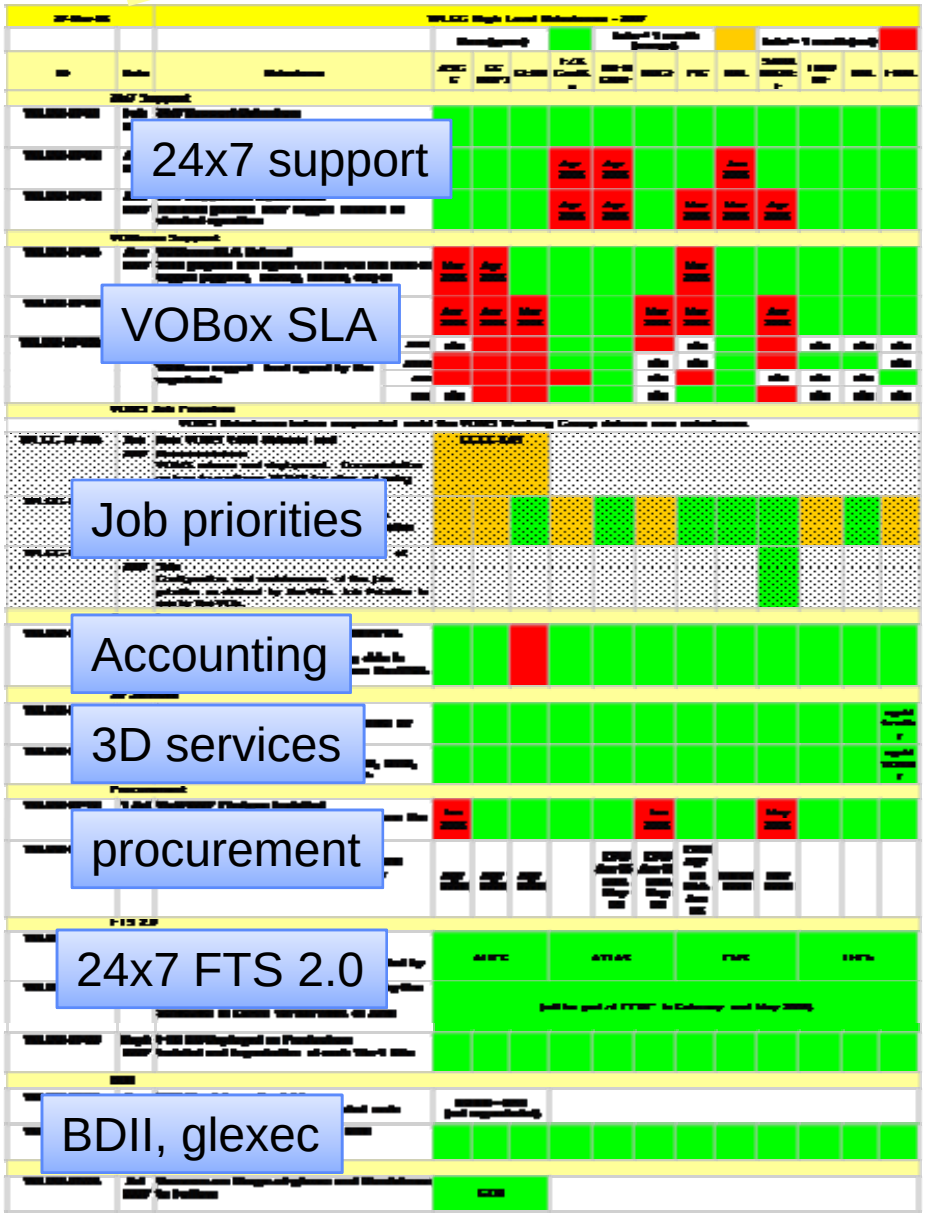

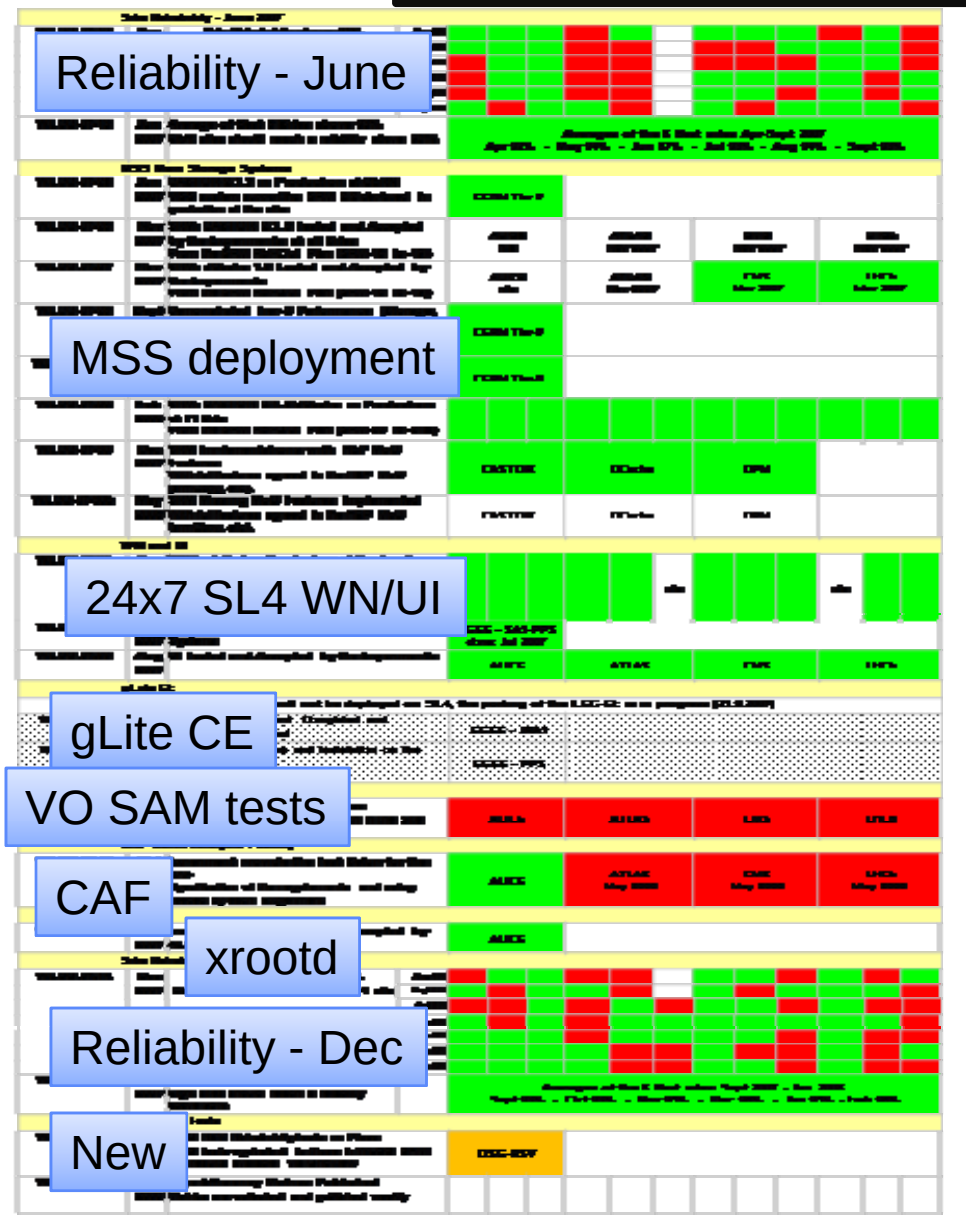

29 Nov 07

Level 1 Milestones

Done	Late <1 mon	Late >1 mon
colour coding		

29-Nov-07		WLCG High Level Milestones - 2007												
			Done (green)				Late < 1 month (orange)				Late > 1 month (red)			
ID	Date	Milestone	ASGC	CC IN2P3	CERN	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF	BNL	FNAL
		24x7 Support												
WLCG-07-01	Feb 2007	**24x7 Support Definition** Definition of the levels of support and rules to follow, depending on the issue/alarm				Sep 2007				Dec 2007				
WLCG-07-02	Apr 2007	**24x7 Support Tested** Support and operation scenarios tested via realistic alarms and situations							Dec 2007	Jan 2008				
WLCG-07-03	Jun 2007	**24x7 Support in Operations** The sites provides 24x7 support to users as standard operations							Feb 2008	Mar 2008				
		VOBoxes Support												
WLCG-07-04	Apr 2007	**VOBoxes SLA Defined** Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes												
WLCG-07-05	May 2007	**VOBoxes SLA Implemented** VOBoxes service implemented at the site according to the SLA						Nov 2007						
WLCG-07-05b	Jul 2007	**VOBoxes Support Accepted by the Experiments** VOBoxes support level agreed by the experiments — ALICE	n/a						n/a			n/a	n/a	n/a
		ATLAS												n/a
		CMS						n/a				n/a	n/a	
		LHCb	n/a					n/a				n/a	n/a	n/a

Only 6 sites have tested their 24 X 7 support, and only 5 have put the support into operation

Only 4 sites have completed the set of VO BOX milestones

27 Mar 08

Level 1 Milestones

Done	Late <1 mon	Late >1 mon

colour coding

27-Mar-08		WLCG High Level Milestones - 2007												
			Done (green)				Late < 1 month (orange)				Late > 1 month (red)			
ID	Date	Milestone	ASGC	CC IN2P3	CERN	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF	BNL	FNAL
		24x7 Support												
WLCG-07-01	Feb 2007	24x7 Support Definition Definition of the levels of support and rules to follow, depe												
WLCG-07-02	Apr 2007	24x7 Suppo Support and realistic alar								an 08				
WLCG-07-03	Jun 2007	24x7 Support in Operations The sites provides 24x7 support to users as standard operations				Apr 2008	Apr 2008		Mar 2008	Mar 2008	Apr 2008			
		VOBoxes Support												
WLCG-07-04	Apr 2007	VOBoxes SLA Defined Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes	Mar 2008	Apr 2008					Mar 2008					
WLCG-07-05	May 2007	VOBoxes SL VOBoxes se according to t							Mar 2008		Apr 2008			
WLCG-07-06b	Jul 2007	VOBoxes Su Experiments VOBoxes support level agreed by the experiments							n/a			n/a	n/a	n/a
		ATLAS						n/a	n/a					n/a
		CMS						n/a			n/a	n/a	n/a	
		LHCb	n/a					n/a				n/a	n/a	n/a

Must be in place for May; understood by Tier 1s now after February experience

Very slow: important to have this defined rapidly now

Only 6→9 sites have tested their 24 X 7 support, and only 5→7 have put the support into operation

Only 4→6 sites have completed the set of VO BOX milestones

SRM v2.2 Deployment

- Deployment plan was defined and agreed last September, but schedule was very tight
- Deployment of dCache 1.8.x and Castor with srm v2.2 was achieved at all Tier0/Tier 1 by December
  - Today 174 srm v2 endpoints are in production
- During February phase of CCRC'08 relatively few problems were found:
  - Short list of SRM v2 issues highlighted, 2 are high priority
  - Will be addressed with fixes or workarounds for May
  - Effort in testing was vital
- Still effort needed in site configurations of MSS – iterative process with experience in Feb & May

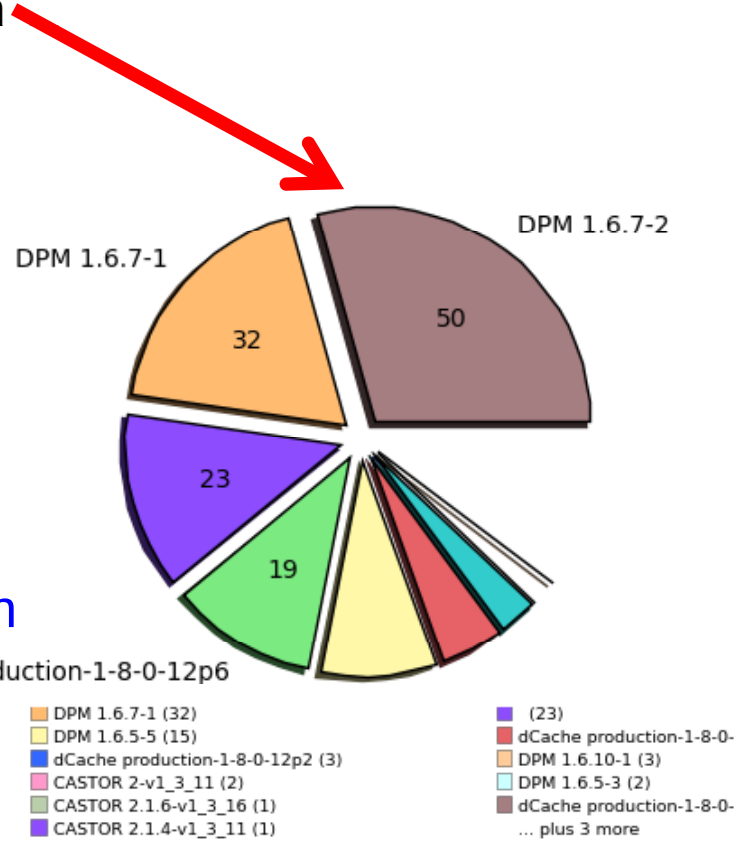

Ian.Bird@cern.ch

Castor performance – Tier 0

Migration policy mod.

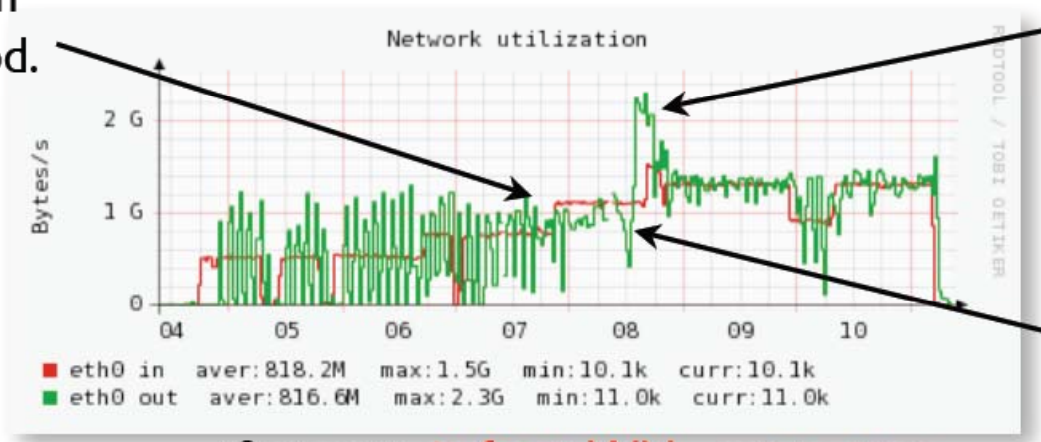

Recovered

Ran out of tapes

t0export: in from WN, out to tape

- CMS:
  - Aggregate rates in/out of castor of 3-4 GB/s
  - Sustained rate to tape 1.3 GB/s with peaks > 2 GB/s
- May:
  - Need to see this with all experiments

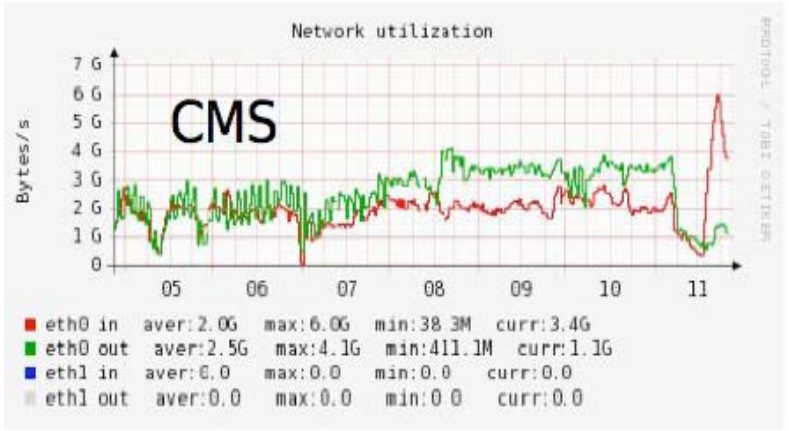

CCRC'08

- Preparation and execution of February phase of CCRC'08 was the most significant activity since December
  - ➔ Full details in Jamie Shiers' talk

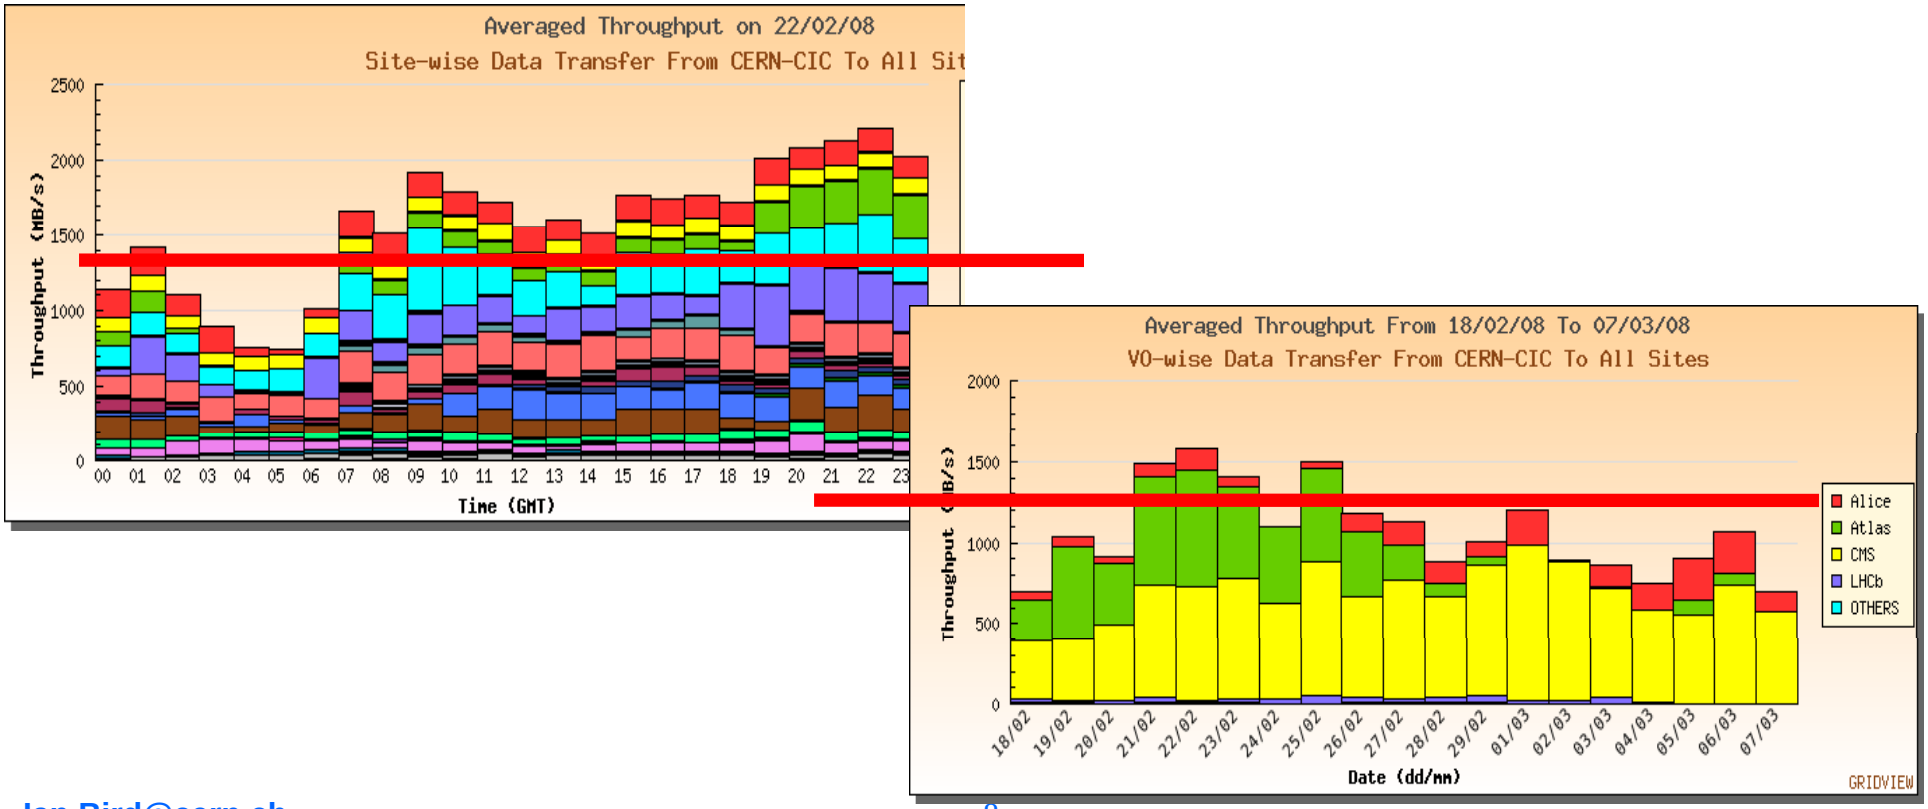

Resource utilization Tier 0+Tier 1

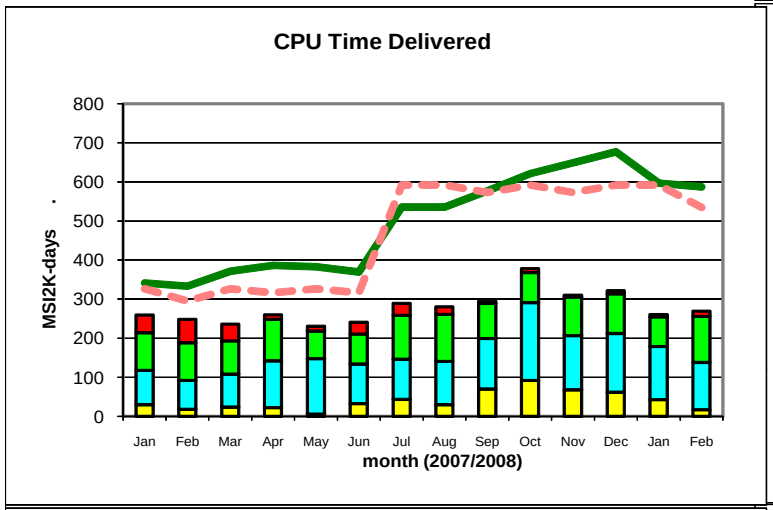

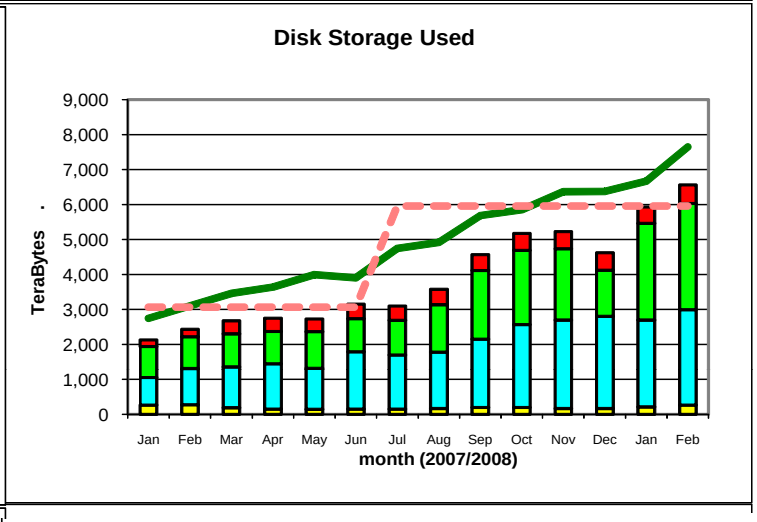

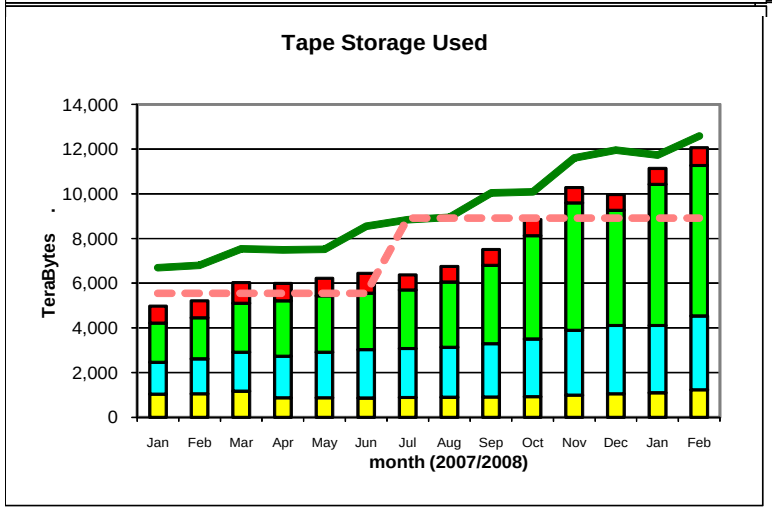

ALICE ATLAS LHCb CMS

installed capacity (inc. efficiency factor)

MoU commitment (inc. efficiency factor)

- Installed capacity exceeds 2007 pledges
  - But – ramp up to 2008 – due in April

Resource Utilization Tier 2

- 52 of 57 federations are reporting
- 114 identified Tier 2 sites
- 107 publish accounting data
  - (was 113/102 in last report

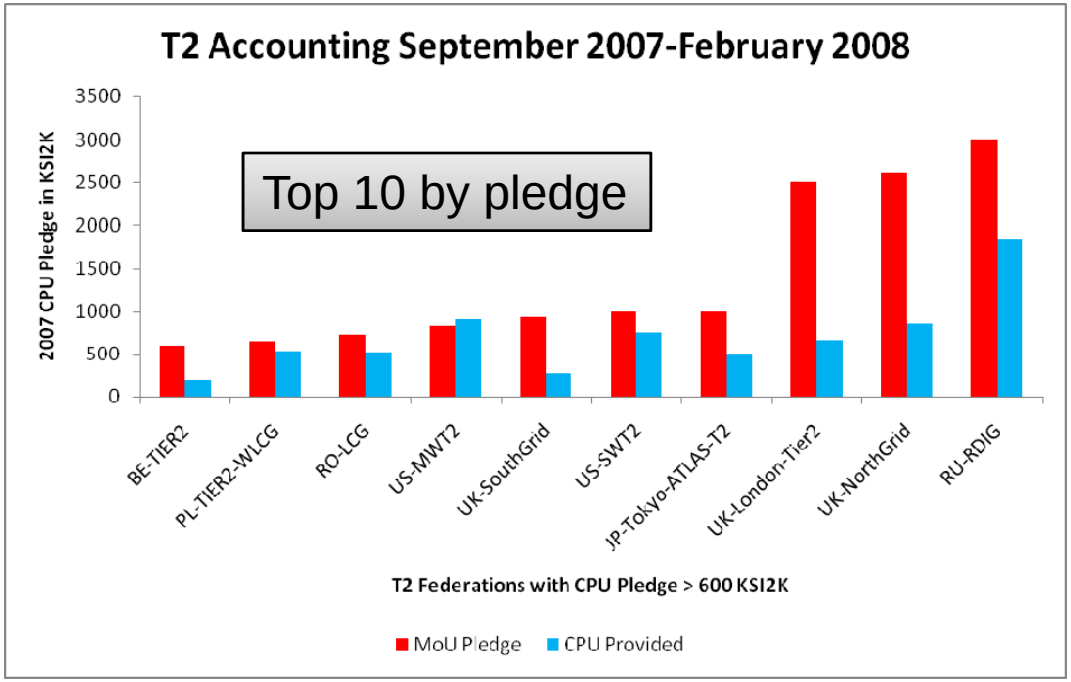

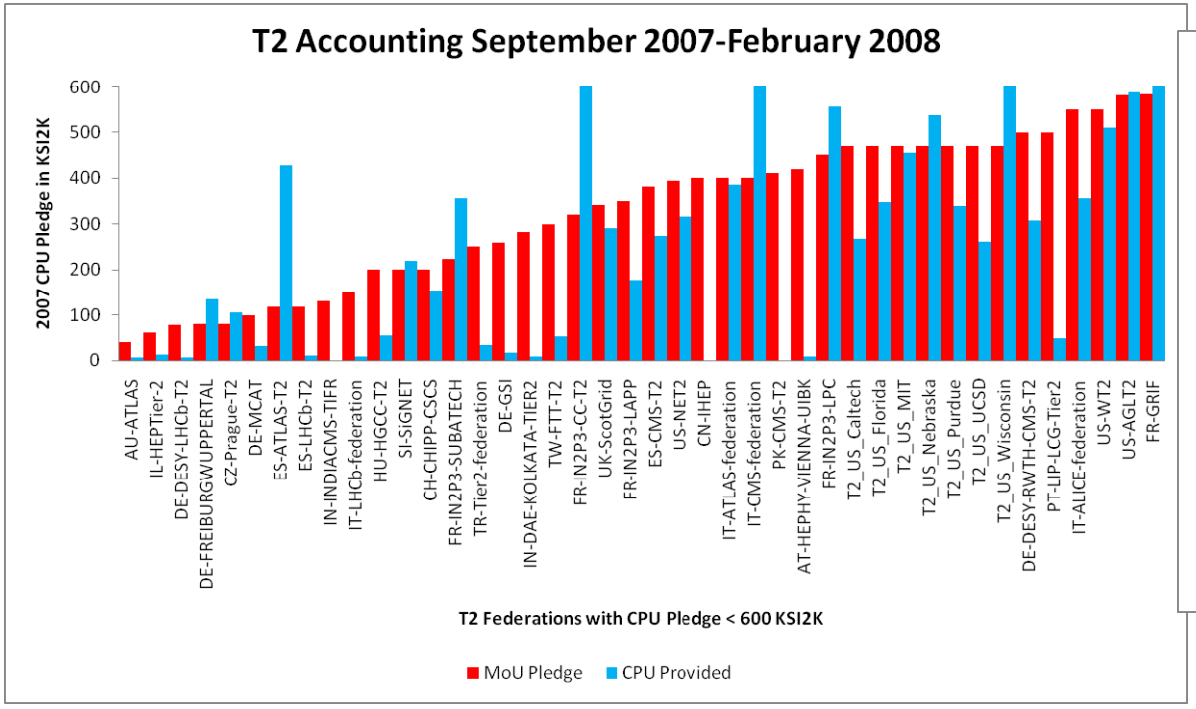

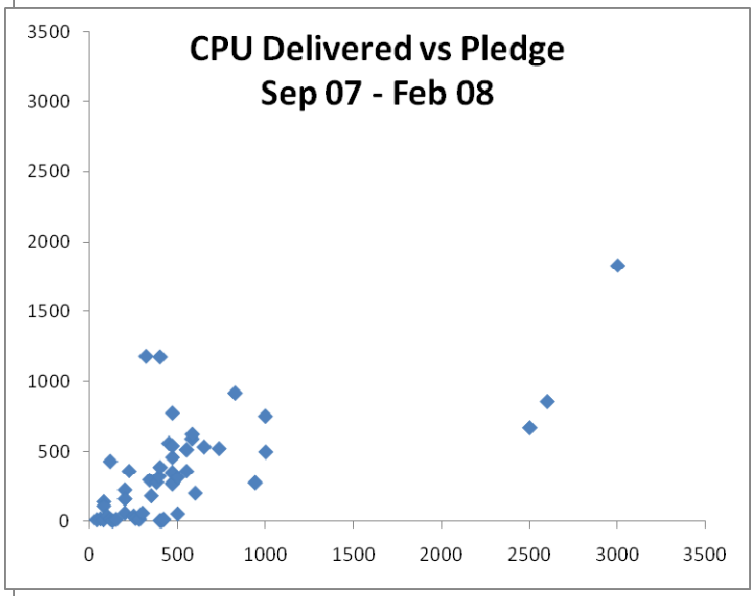

Tier 2 CPU: Sep 07 – Feb 08

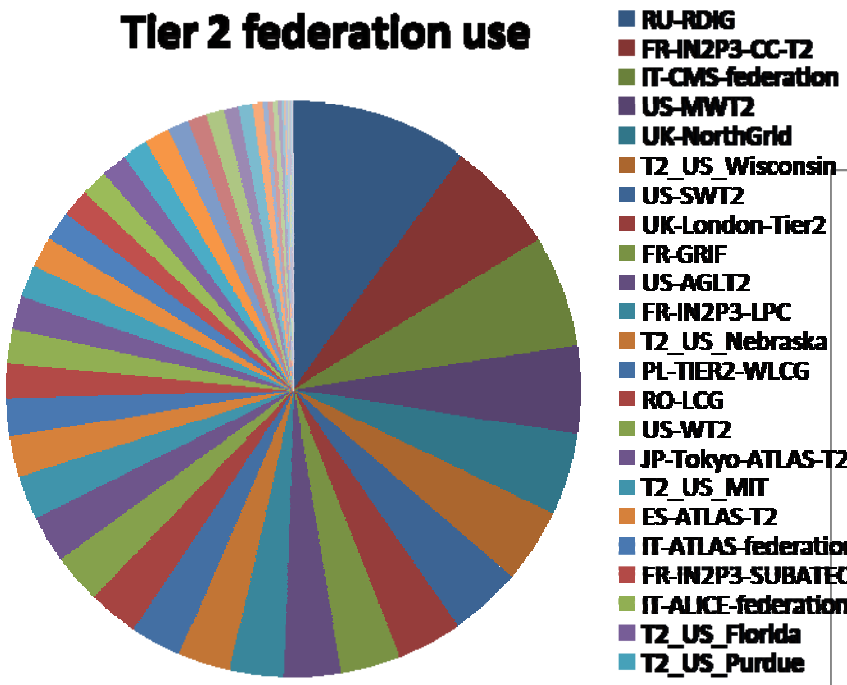

Missing federations:

Finland
India (IN-INDIACMS-TIFR)
Norway
Sweden
Ukraine

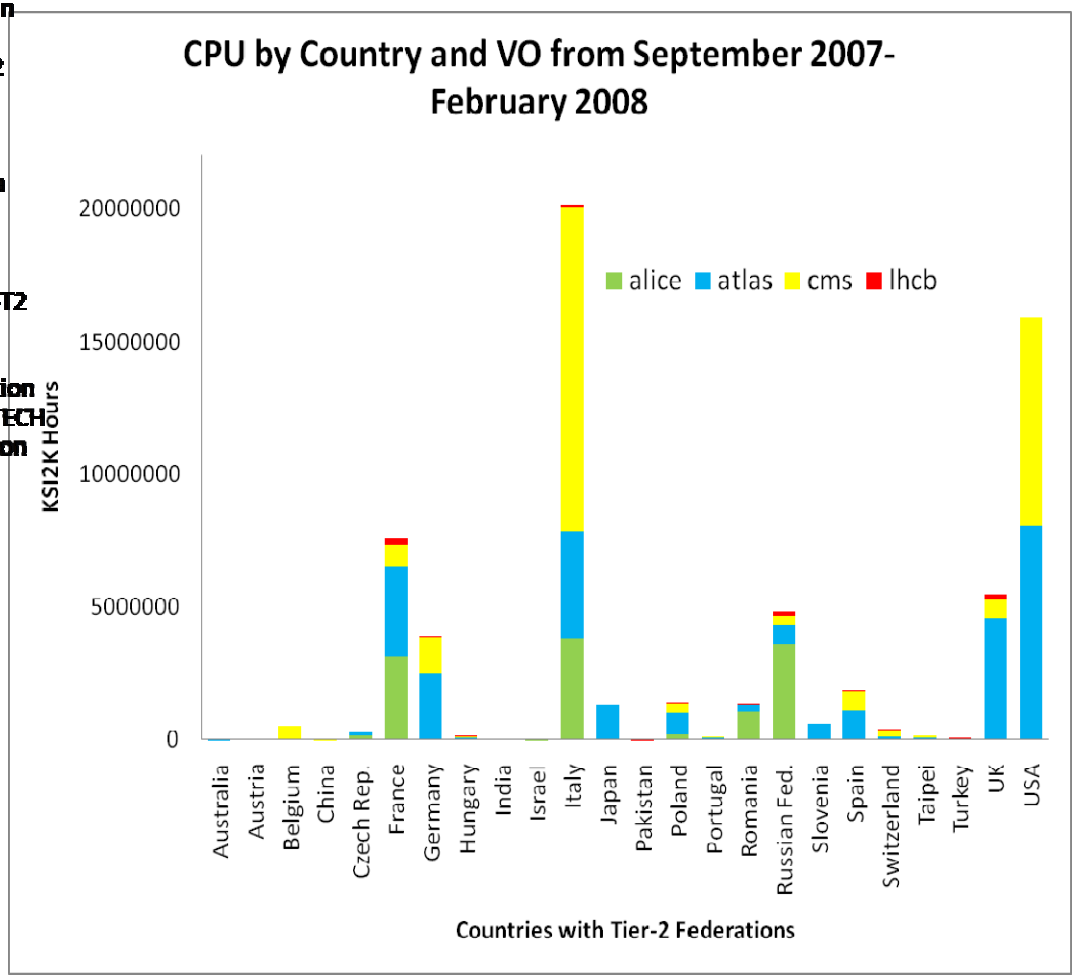

Ian.Bird@cern.ch

Workloads

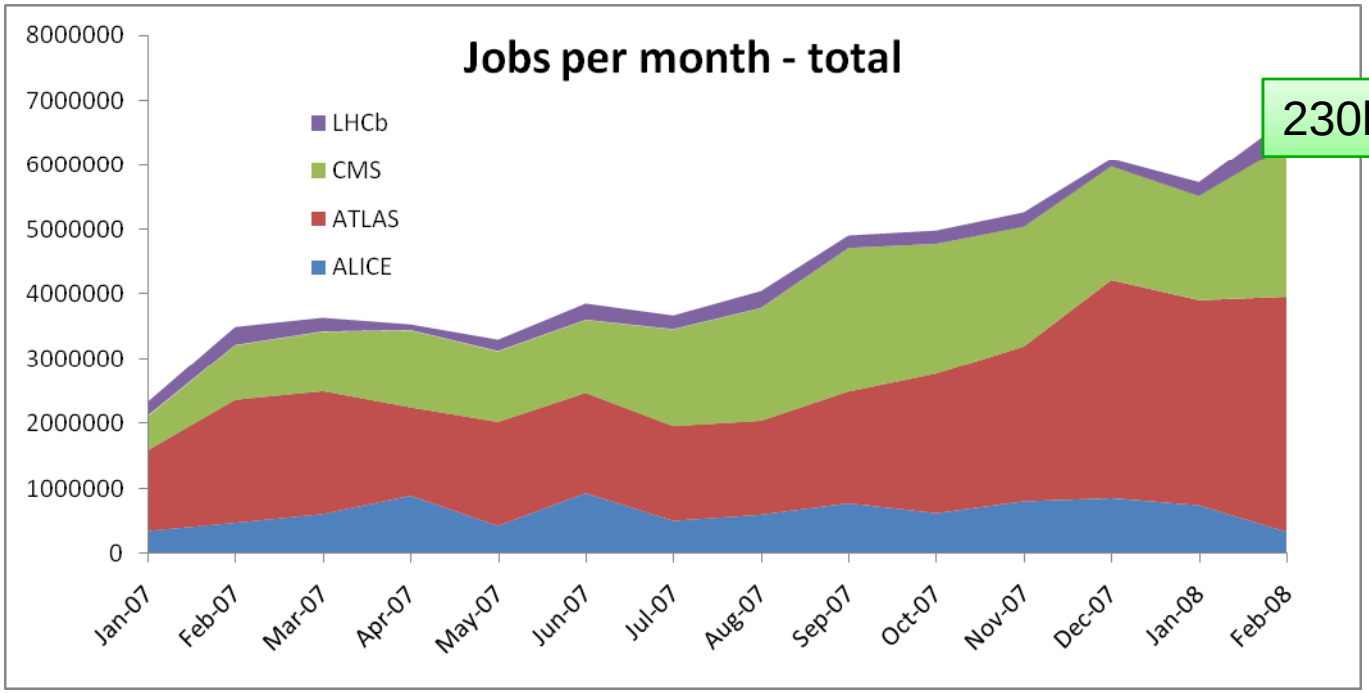

230k /day

- These workloads are at the level anticipated for 2008 data taking

Resource ramp up for 2008

- CPU:
  - Most Tier 1 sites will have full 2008 pledges in place for 1 May
  - Total of 36725 KSi2K.
  - Largest missing is +2500 at NL-T1 due Nov.
- Disk and tape
  - Many sites will catch up later in the year as need expands:
  - 2008 disk requirements are 23 PB with15.5 PB expected by 1 May
  - 2008 tape requirements are 24 PB with 15 PB expected by1 May.
- May run of CCRC'08 at 55% only requires +1PB of disk and +1.5PB of tape (mostly reusable) so should have no resource problems.
- Full status of resource installation will be reported at C-RRB in April.
- **Many sites had problems with procurement process/ vendor delivery/ faulty equipment**
  - ➢ **These issues must be taken into account in future – the process is long, but yearly deadlines are important**

Resource pledges vs requirements

Tier 1	ALICE	ATLAS	CMS	LHCb	Sum 2008
CPU	-45%	6%	7%	43%	-5%
Disk	-40%	2%	-23%	33%	-12%
Tape	-49%	-5%	-4%	39%	-13%
Tier 2					
CPU	-46%	0%	27%	-7%	-3%
Disk	-20%	-19%	-16%	1443%	-15%

Situation as of 26/3/08

Tier 1	2008	2009	2010	2011	2012
CPU	-5%	-11%	-11%	-17%	-24%
Disk	-12%	-12%	-15%	-17%	-24%
Tape	-13%	-13%	-17%	-22%	-29%
Tier 2					
CPU	-3%	-14%	-34%	-37%	-43%
Disk	-15%	-4%	-1%	-11%	-21%

Tier 0/Tier 1 Site reliability

- Target:
  - Sites 91% & 93% from December
  - 8 best: 93% and 95% from December
- See QR for full status

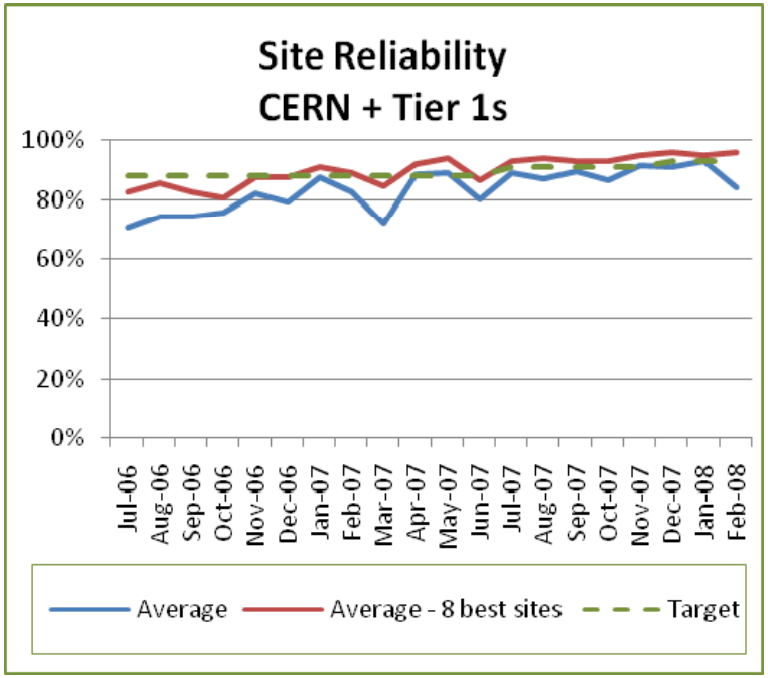

	Sep 07	Oct 07	Nov 07	Dec 07	Jan 08	Feb 08
All	89%	86%	92%	87%	89%	84%
8 best	93%	93%	95%	95%	95%	96%
Above target (+>90% target)	7 + 2	5 + 4	9 + 2	6 + 4	7 + 3	7 + 3

Tier 2 Reliabilities

- Reliabilities published regularly since October

Overall	Top 50%	Top 20%	Sites
76%	95%	99%	89→100

- In February 47 sites had > 90% reliability

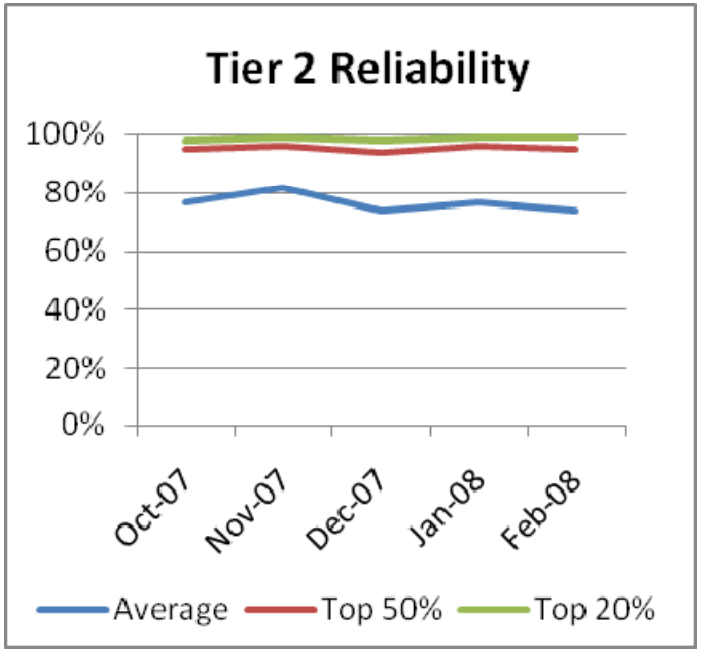

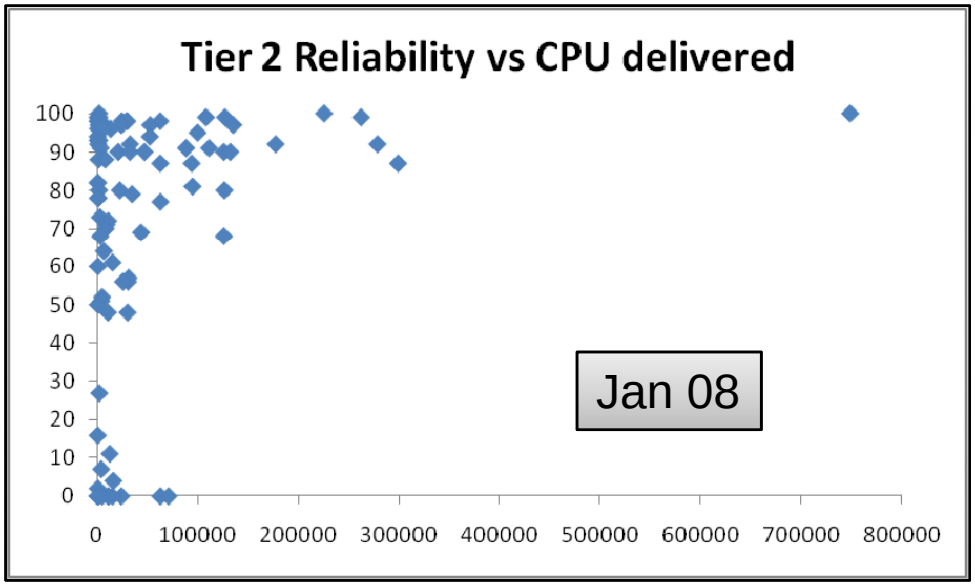

- For the Tier 2 sites reporting:

Sites	Top 50%	Top 20%	Sites> 90%
%CPU	72%	40%	70%

- For Tier 2 sites not reporting, 12 are in top 20 for CPU delivered

Reliability reporting

- Currently (Feb 08) All Tier 1 and 100 Tier 2 sites report reliabilities
- Recent progress: MB set up group to
  - Agreement on equivalence of NDGF tests with those used at EGEE and all other Tier 1 sites – now in production at NDGF
    - Should also be used for Nordic Tier 2 sites
  - Similar process with OSG (for US Tier 2 sites): tests only for CE so far, agreement on equivalence, tests are in production, publication to SAM in progress
    - Missing – SE/SRM testing
    - Expect full production May 2008 (new milestone introduced)
- Important that we have all Tier 2s regularly tested and reporting
- VO-specific reliabilities should be a focus now

Progress in EGEE-III

- EGEE-III now approved
  - Starts 1st May, 24 months duration (EGEE-II extended 1 month)
- Objectives:
  - Support and expansion of production infrastructure
  - Preparation and planning for transition to EGI/NGI
- Many WLCG partners benefit from EGEE funding, especially for grid operations: effective staffing level is 20-25% less
  - Many tools: accounting, reliability, operations management funded via EGEE
  - Important to plan on long term evolution of this
- Funding for middleware development significantly reduced
- Funding for specific application support (inc HEP) reduced

➢ Important for WLCG that we are able to rely on EGEE priority on operations, management, scalability, reliability

Comments on EGI design study

- Goal is to have a fairly complete blueprint in June
- Main functions presented to NGIs in Rome workshop in March
- Essential for WLCG that EGI/NGI continue to provide support for the production infrastructure after EGEE-III
  - We need to see a clear transition and assurance of appropriate levels of support
  - Transition will be 2009-2010
    - Exactly the time that LHC services should not be disrupted
- Concerns:
  - NGIs did not all seem convinced that a large European production-quality infrastructure is a goal
    - Not clear that there is agreement on the scope – hard to see how a clear agreed blueprint will be available in June
  - Tier 1 sites (and existing EGEE expertise) not well represented by many NGIs
- WLCG representatives must approach their NGI reps and ensure that EGI/NGIs provide the support we need

Evolution of capacity Tier 0/CAF

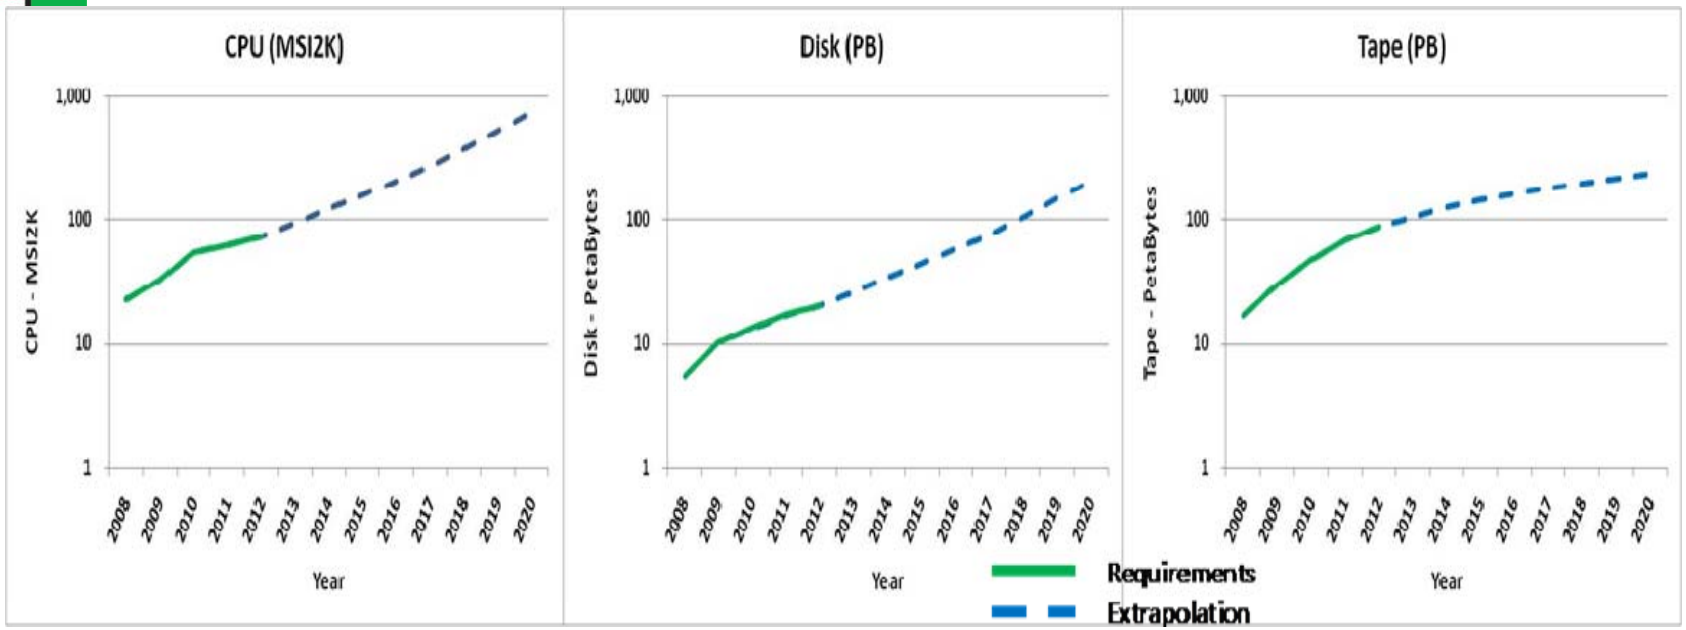

- Evolution – assumes total cost within budget capped at level at the end of current medium term plan
  - Result is average annual rate of increase in CPU and disk of 30% after 2012
  - Tape capacity linear growth at 19 PB/year
- This is low compared to experience in last 15 years – 100%/year
- Experience shows that needs are usually underestimated ...
- Use this profile for planning – but supports only the basic physics programme
  - Operations budget must fund the computing infrastructure as well as the equipment

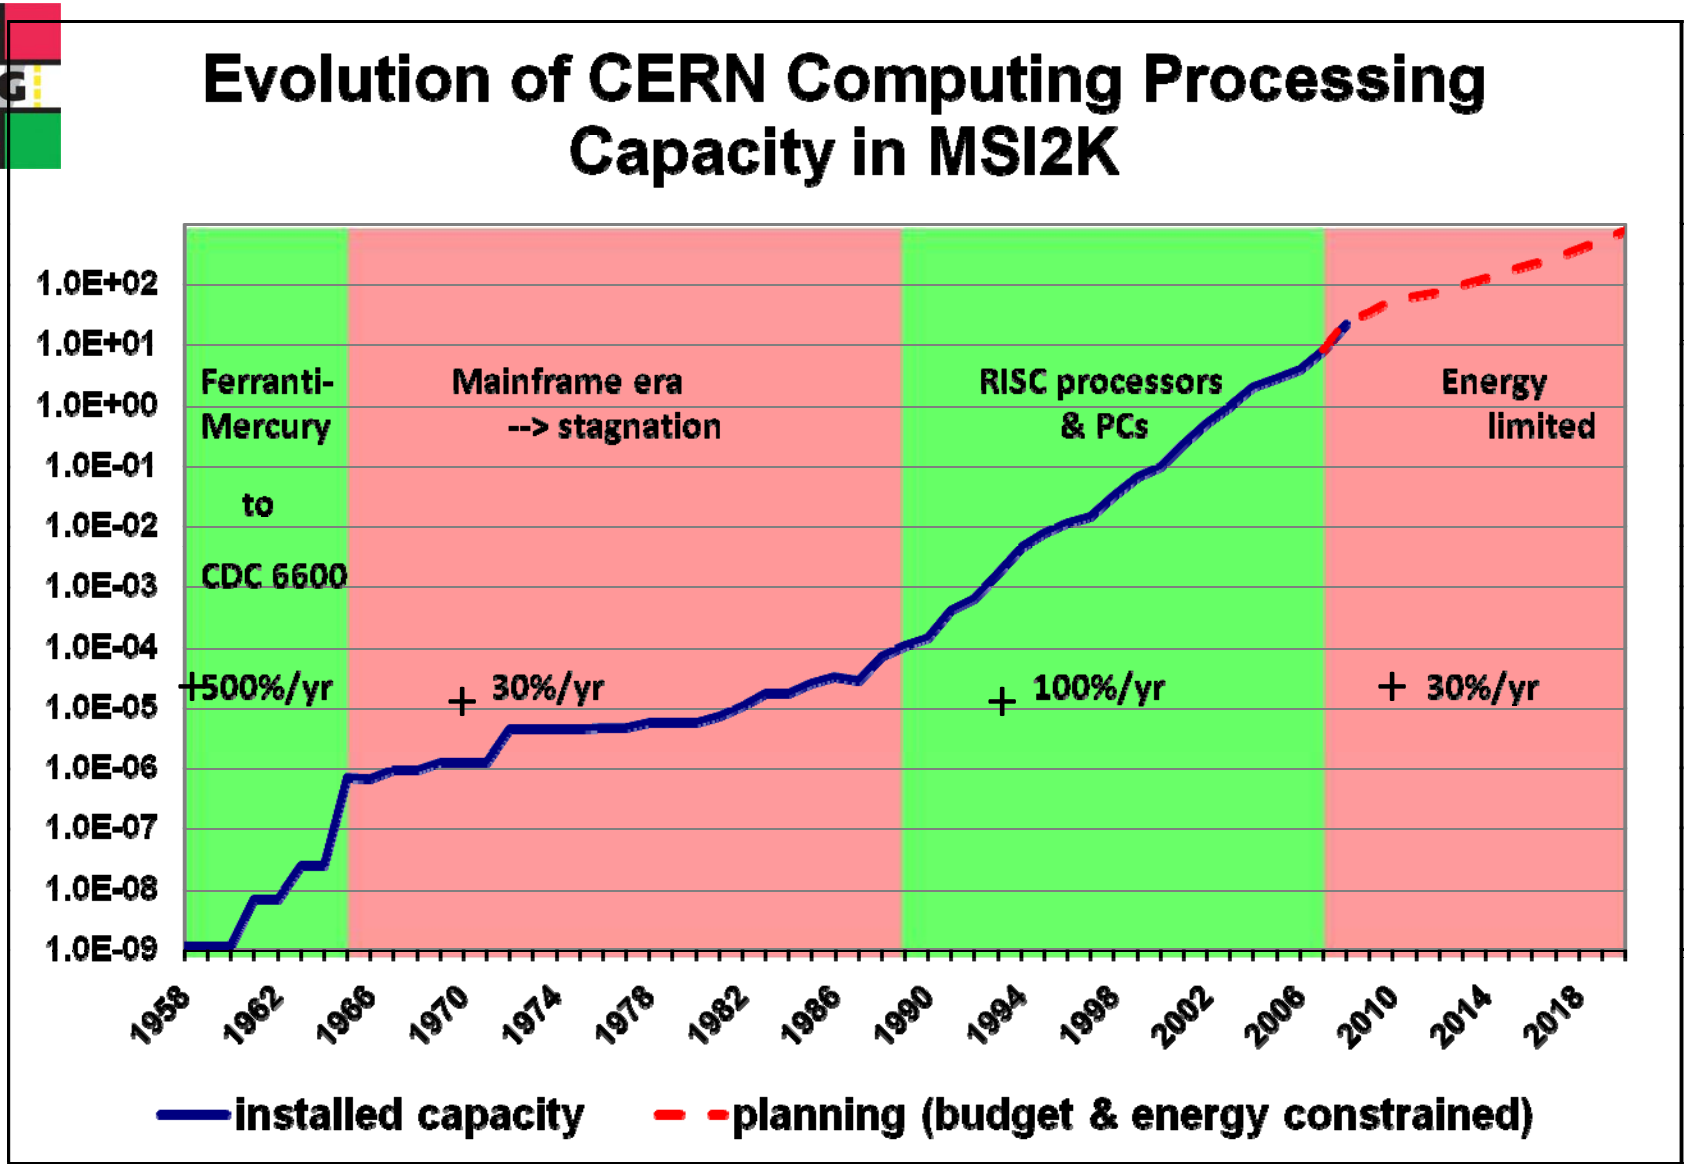
Evolution of CERN Computing Processing Capacity in MSI2K
1.0E+02
1.0E+01
1.0E+00
1.0E-01
1.0E-02
1.0E-03
1.0E-04
1.0E-05
1.0E-06
1.0E-07
1.0E-08
1.0E-09
Ferranti-
Mercury
to
CDC 6600
Mainframe era
--> stagnation
RISC processors
& PCs
Energy
limited
+500%/yr
+ 30%/yr
+ 100%/yr
+ 30%/yr
1958
1962
1966
1970
1974
1978
1982
1986
1990
1994
1998
2002
2006
2010
2014
2018
installed capacity
planning (budget & energy constrained)

Power requirements

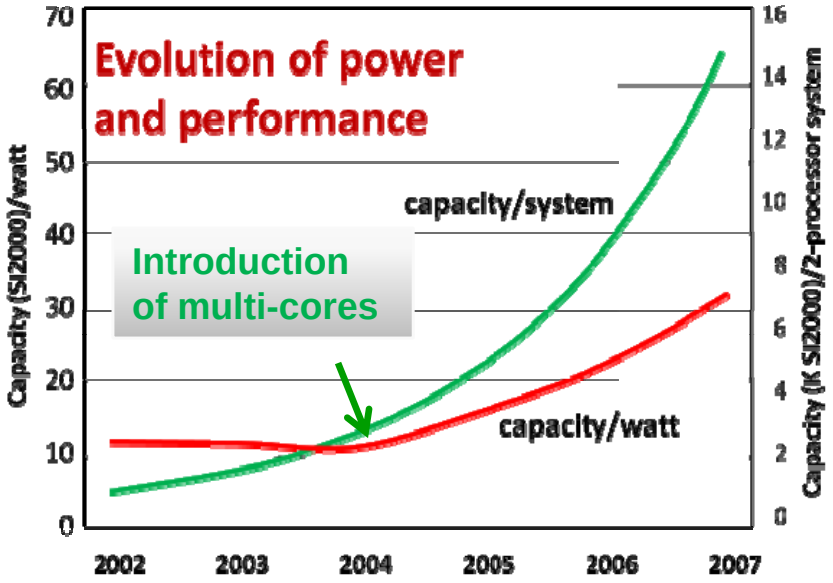

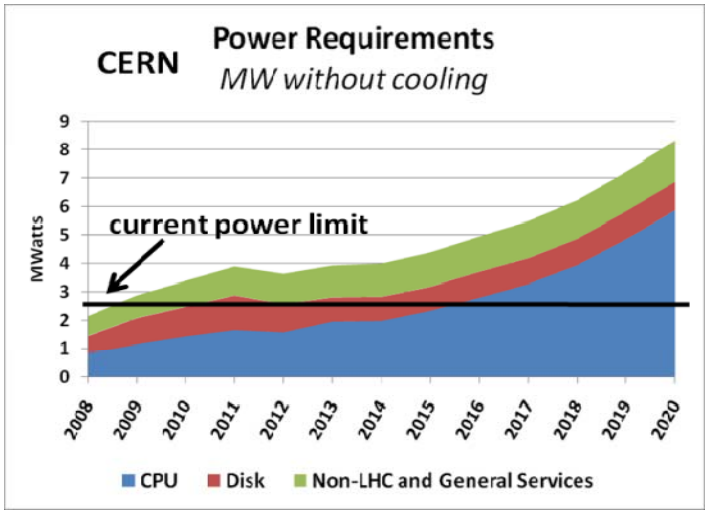

- Expect power requirements to grow with capacity of CPU
  - This is not a smooth process: depends on new approaches and market-driven strategies (hard to predict) e.g. improvement in cores/chip is slowing; power supplies etc. already >90% efficient
  - No expectation to get back to earlier capacity/power growth rate

- Existing CERN Computer Centre will run out of power in 2010
  - Current usable capacity is 2.5MW
  - Situation will only get worse

Strategies

- Major investments are needed for new Computer Centre infrastructure at CERN and major Tier 1 centres
  - IN2P3, RAL, FNAL, BNL, SLAC already have plans
  - IHEPCCC report to ICFA at DESY in Feb '08
- At CERN this can be done within the agreed Physics Computing budget as long as it remains constant at 2012 levels – this is the conclusion of study over the last year
  - Now have technical info needed to start a formal acquisition process
- Estimated time ranges from 27 – 43 months
- Even if we start now will also need short term strategies
  - e.g. External hosting – expensive

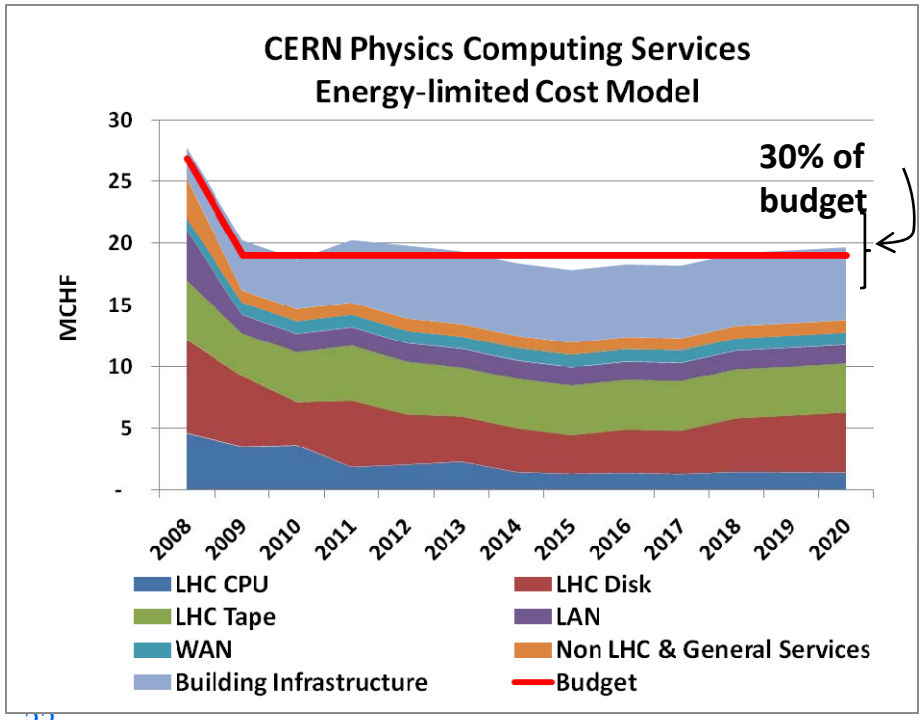

Process

- Currently no authorisation to proceed with the formal process ...
- Consequences
  - Growth in capacity for Tier 0 and CAF after 2009 will be limited to what can be achieved by replacement of existing equipment with newer, more efficient, systems
  - Difference between capacity at CERN and Tier 0/CAF requirements would need to be made up by increases at Tier 1s, or experiments must reconsider their computing strategies

- C-RRB needs to be made aware of this issue – CERN will not be able to fulfil its commitments to LHC computing beyond 2009

Summary

- CCRC'08 phase 1 has been a success – but significant work still needed for May and data taking preparation
  - SRM v2.2 deployment was done, but work is still needed on configuration of the MSS systems together with experiments
  - Tuning of tape access with real use patterns – may require experiments to reconsider analysis patterns
- Resource ramp-up: based on experiences and problems with 2008 procurements
  - Must ensure in future years that allowance is made for delays and problems
  - Important that the yearly April schedules are met – to be ready for accelerator start ups
- Remaining Tier 2 federations must now ensure that they regularly report (and verify) accounting and reliability data

Summary

- WLCG – especially Tier 1s – should influence the directions of the EGI Design study
  - Must ensure that we see a clear and appropriate strategy emerging that is fully supported by the NGIs
  - Must engage the NGI representatives in this
- Current CERN Computer Centre power limitations will have consequences for future computing capacity growth
  - Will have to be brought to the attention of the C-RRB
  - Important that there is a clear plan to address this