


Status of the WLCG Tier-2 Centres

M.C. Vetterli
Simon Fraser University
and TRIUMF


LHCC Review,
CERN, July 1st 2008



Status of the WLCG Tier-2 Centres

OUTLINE

1. *Tier-2 Roles*
2. *Tier-2 Sites & Capacities*
3. *Tier-2 Performance:
CCRC'08 & otherwise*
4. *Tier-2 Issues as we
approach LHC startup*



Acknowledgements



- *This is really a rapporteur talk; many people to thank for input*
- *Experiment computing coordinators: F. Carminati (ALICE), D. Barberis (ATLAS), M. Kasemann (CMS), N. Brook (LHCb)*
- *From CCRC'08 post-mortem workshop: L. Betev (ALICE), S. Campana (ATLAS), D. Bonacorsi (CMS), N. Brook (LHCb)*
- *WLCG: I. Bird, J. Gordon, J. Shiers, M. Jouvin, J. Salt, H. Hoorani*
- *Others: K. Bloom (US-CMS), M. Ernst (US-ATLAS), N. Geddes (UK), J. Coles (UK), ...*

Roles of the Tier-2 Centres

- *Monte Carlo production centres; data uploaded to Tier1s*
→ *well tested for a long time*
- *User analysis, mostly from AODs, although physicists will also get access to ESDs and RAW data at the Tier2s*
→ *chaotic analysis less well tested*
- *Exception for LHCb where T2s only do Monte Carlo*
- *ATLAS uses the cloud model and CMS has T2s associated to T1s. However for ALICE, the T2s access all T1s equally*

Tier-2 Sites & Capacities

- A little difficult to determine exact resources installed in the Tier2s;
 → must be improved



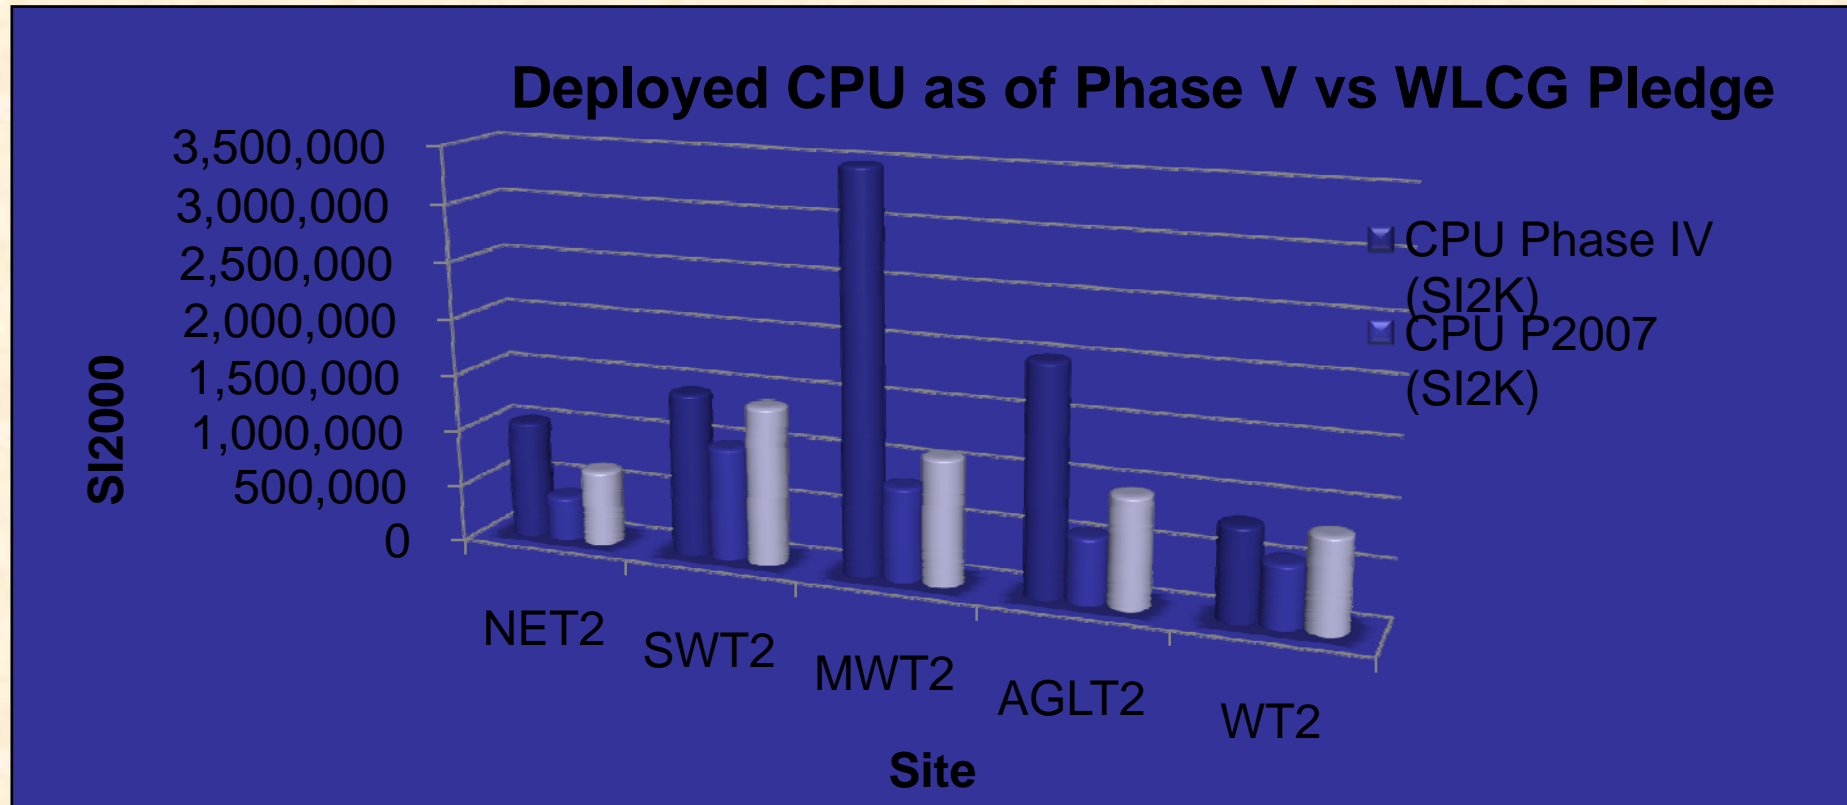
(Ian Bird, March LCG OB)

- But is provided < pledged because of availability or lack of use?
 → need better reporting


Tier-2 Sites & Capacities

- *At the time of the March OB, Ian had “114 identified Tier-2 sites”, but not all sites were reporting.*
 - *I get 136 sites from the MoU*
 - *Pledges for 2008 add up to over 46 MSI2k of CPU and just under 12.5 PB of disk across all Tier-2s*
 - *But how much is actually installed?*
 - e.g. US-CMS: 7 sites pledged 0.9 MSI2k and 200 TB each*
 - 4 sites already there, 2 OK for either CPU or disk*
 - All sites complete by the end of the summer*
- Probably not as good everywhere*

US-ATLAS Tier-2 CPU Capacities



US-ATLAS Tier-2 Disk Capacities



≥ 100 TB being added to each site soon

Status of Tier-2 Pledges



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%



Simon Fraser

ian.bird@cern.ch

14




Tier-2 Performance Summary

- *Overall, the Tier-2s are contributing much more now*
- *Significant fractions of the Monte Carlo simulations are being done in the T2s for all experiments*
- *Reliability is better, but still needs to improve*
- *CCRC'08 exercise is generally considered a success for the Tier2s*

LHCb Monte Carlo Production

Total Running Jobs: 402
DIRAC: 11.44% LCG: 87.81%




39 sites participating, most of them Tier2s



Simon Fraser



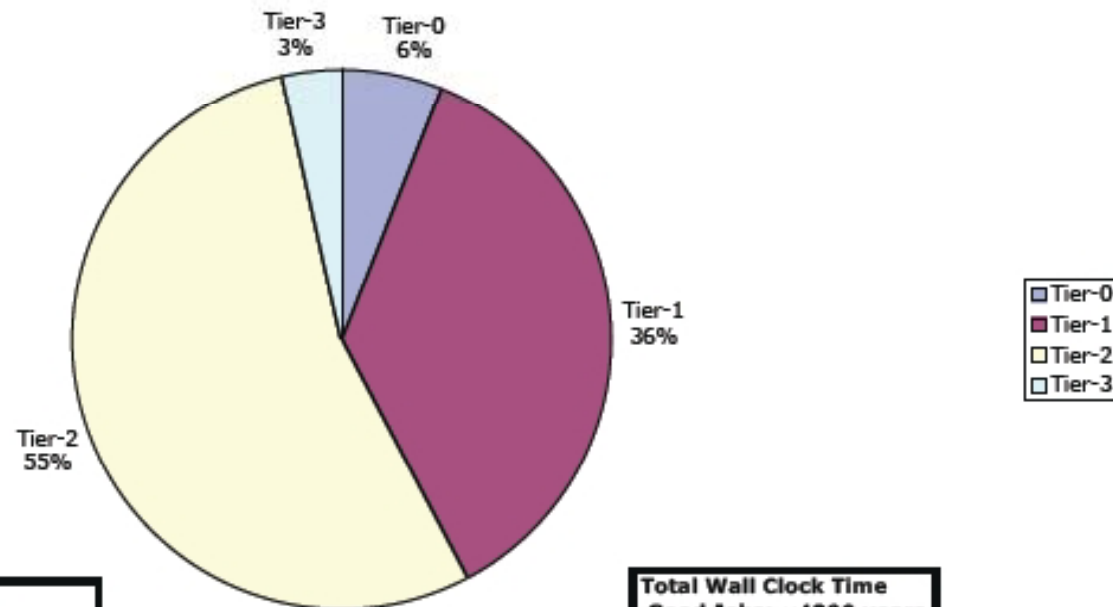
ALICE Monte Carlo Production



ATLAS Monte Carlo Production

≈ 55% of production from Tier-2s

ATLAS Production in 2007



Source:
ATLAS Production Database
Wall Clock Time

Total Wall Clock Time
Good Jobs: ~4300 years
Bad Jobs: ~ 700 years

ATLAS Week 14 Feb., 2008

17



Simon Fraser

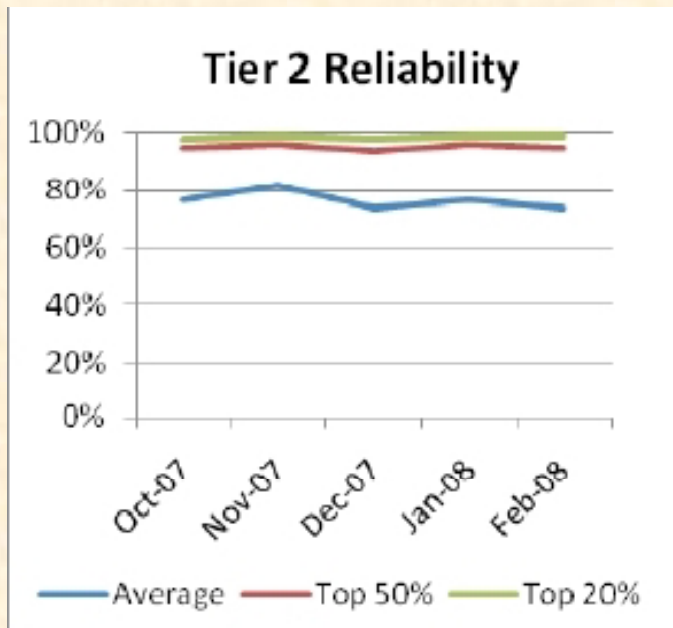


Tier-2 Reliability

A lot of green (>90%),

but still too much yellow (60-90%),

and too many sites N/A (mostly OSG and NDG)



Tier-2 Availability and Reliability Report

Federation Summary - Sorted by Reliability

May 2008

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

Availability = % of successful tests
 Reliability = Availability / Scheduled Availability
 Reliability and Availability for federation - average of all sites in the federation

Colour coding : N/A < 30% < 60% < 90% >= 90%

Federation	Reliability	Availability	Federation	Reliability	Availability
FR-IN2P3-LAPP	100 %	100 %	TR-Tier2-federation	82 %	82 %
FR-GRIF	99 %	99 %	EE-NICPB	80 %	70 %
AT-HEPHY-VIENNA-UIBK	99 %	94 %	DE-FREIBURGWUPPERTAL	73 %	63 %
DE-DESY-ATLAS-T2	99 %	98 %	DE-MCAT	72 %	64 %
JP-Tokyo-ATLAS-T2	98 %	97 %	HU-HGCC-T2	70 %	63 %
FR-IN2P3-LPC	98 %	98 %	US-NET2	N/A	N/A
TW-FTT-T2	98 %	98 %	US-MWT2	N/A	N/A
FR-IN2P3-CC-T2	98 %	98 %	DE-DESY-RWTH-CMS-T2	66 %	66 %
US-SWT2	N/A	N/A	IN-INDIACMS-TIFR	62 %	54 %
SI-SIGNET	96 %	96 %	PK-CMS-T2	62 %	60 %
FR-IN2P3-SUBATECH	96 %	96 %	IN-DAE-KOLKATA-TIER2	61 %	57 %
ES-CMS-T2	95 %	93 %	KR-KISTI-T2	59 %	59 %
CH-CHIPP-CSCS	94 %	94 %	US-AGLT2	N/A	N/A
UK-London-Tier2	94 %	73 %	IL-HEPTier-2	43 %	43 %
UK-NorthGrid	93 %	93 %	AU-ATLAS	20 %	20 %
ES-ATLAS-T2	93 %	90 %	DE-GSI	N/A	N/A
UK-ScotGrid	92 %	75 %	FI-HIP-T2	N/A	N/A
PL-TIER2-WLCG	92 %	90 %	NO-NORDGRID-T2	N/A	N/A
IT-ALICE-federation	91 %	87 %	SE-SNIC-T2	N/A	N/A
IT-ATLAS-federation	91 %	87 %	T2_US_Caltech	N/A	N/A
IT-CMS-federation	91 %	87 %	T2_US_Florida	N/A	N/A
IT-LHCb-federation	91 %	87 %	T2_US_MIT	N/A	N/A
CA-EAST-T2	90 %	90 %	T2_US_Nebraska	N/A	N/A
CZ-Prague-T2	89 %	79 %	T2_US_Purdue	N/A	N/A
UK-SouthGrid	88 %	85 %	T2_US_UCSD	N/A	N/A
CN-IHEP	85 %	84 %	T2_US_Wisconsin	N/A	N/A
RO-LCG	84 %	78 %	UA-	N/A	N/A
PT-LIP-LCG-Tier2	84 %	77 %	US-WT2	N/A	N/A
CA-WEST-T2	84 %	83 %			
ES-LHCb-T2	83 %	83 %			
BE-TIER2	83 %	82 %			
RU-RDIG	82 %	81 %			



Simon Fraser



Tier-2 Monitoring in the UK

UK Grid Test Results

http://pprc.qmul.ac.uk/~lloyd/gridpp/uktest.html

Getting Started Latest Headlines

Most recent job submitted on Mon Jun 30 2008 at 09:50

ATLAS Analysis jobs (Release 14.2.0) are submitted every 10 mins to any UK CE. Z->e⁺e⁻ data are read from the local SE and analysed. The key is S: Success, A: Aborted, C: Current, F: Failed, X: Cancelled. The second table shows where the jobs ran.

Latest Jobs										Time of Last Job		Overall			
10	9	8	7	6	5	4	3	2	1	Submitted	Successful	24 Hrs	Week	Month	6 Mon
F	F	S	S	S	S	S	S	S	C	30/06/08 09:50	30/06/08 09:30	94%	94%	95%	95%

Total jobs in last week: 985 Successful: 931 Failed: 54 [History](#)

Institute	Jobs Run	Successful	Failed	Recent Failures			
1 RAL PPD	224	23%	210	94%	14	6%	10 9 8 7 6 5 4 3 2 1
2 Glasgow	159	16%	155	97%	4	3%	4 3 2 1
3 RAL Tier-1	117	12%	106	91%	11	9%	10 9 8 7 6 5 4 3 2 1
4 Durham	101	10%	94	93%	7	7%	7 6 5 4 3 2 1
5 Oxford	77	8%	77	100%	0	0%	
6 Birmingham	72	7%	71	99%	1	1%	1
7 Imperial HEP	38	4%	31	82%	7	18%	7 6 5 4 3 2 1
8 Liverpool	36	4%	35	97%	1	3%	1
9 Brunel	34	3%	33	97%	1	3%	1
10 QMUL	33	3%	33	100%	0	0%	
11 RHUL	31	3%	30	97%	1	3%	1
12 Bristol	25	3%	24	96%	1	4%	1
13 Lancaster	23	2%	20	87%	3	13%	3 2 1
14 Cambridge	15	2%	12	80%	3	20%	3 2 1
Tier	Jobs Run	Successful	Failed				
1 SouthGrid	413	42%	394	95%	19	5%	
2 ScotGrid	260	26%	249	96%	11	4%	
3 LondonGrid	136	14%	127	93%	9	7%	
4 Tier-1	117	12%	106	91%	11	9%	
5 NorthGrid	59	6%	55	93%	4	7%	

Tier-2 Centres in CCRC'08 - General

- *Overall, the Tier-2s and the experiments considered the CCRC'08 exercise to be a success*
- *The networking/data transfers were tested extensively; some FTS tuning was needed, but it worked out*
- *Experiments tended to continue other activities in parallel which is a good test of the system, although the load was not as high as anticipated*
- *While CMS did include significant user analysis activities, the chaotic use of the Grid by a large number of inexperienced people is still to be tested*

Tier-2 Centres in CCRC'08 – LHCb/ALICE

- *LHCb concentrated on the T0 & T1 layers; Tier-2s are mainly for Monte Carlo production for LHCb*
- *ALICE upgraded its Grid services and concentrated more on the T0, T1, and CAF layers*


However, they did replicate ESDs to the Tier-2 centres that have large detector communities, since these will be the most active with first data

Tier-2 layer will be exercised more extensively this summer

Tier-2 Centres in CCRC'08 - ATLAS

- *ATLAS concentrated more on data distribution than user analysis tests*
- *However, the full chain was exercised: T0-T1, T1-T1, T1-T2*
- *In fact, transfers were at a higher rate than needed for '08*
- *Some problems with “double registration”; files replicated correctly, but then it is done again for some reason*
- *T1->T2: a complete copy of the AODs at T1 should be replicated at among the T2s, within 6 hours from the end of the exercise*
→ every cloud met this goal!

Tier-2 Centres in CCRC'08 - ATLAS



SIGNET: ATLAS DDM configuration issue (LFC vs RLS)

CSTCDIE: joined very late. Prototype

“suspect”: double registration problem



Simon Fraser



Tier-2 Centres in CCRC'08 - CMS

- *Thoroughly tested data transfers, both within a T1's region and to T2s in other regions*
- *Establishing criteria to “commission” a link*
- *Did extensive testing of user analysis at the Tier-2s*
- *Central and local control of job submission*

Tier-2 Centres in CCRC'08 - CMS

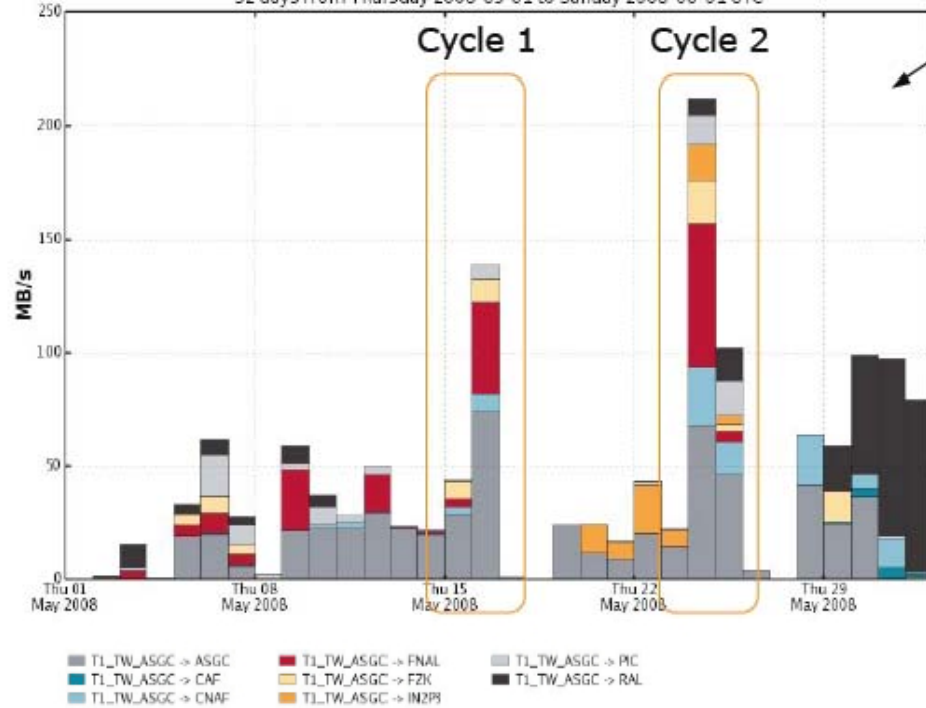
ASGC → T2's

T1->/Region	ASGC	CERN T1	CNAF	FNAL	PZK	IN2P3	PIC	RAL
ASGC	36.8	2.6	8.0	17.7	5.9	8.3	3.4	6.7

regional (pointing to ASGC) non-regional (pointing to FNAL, PZK, IN2P3, PIC)

Daily CMS PhEDEx transfer rate, Debug + Production

site to region links for non-tape storage only
Sources matching 'T1_TW*' and destinations matching 'T2_*'
32 days from Thursday 2008-05-01 to Sunday 2008-06-01 UTC



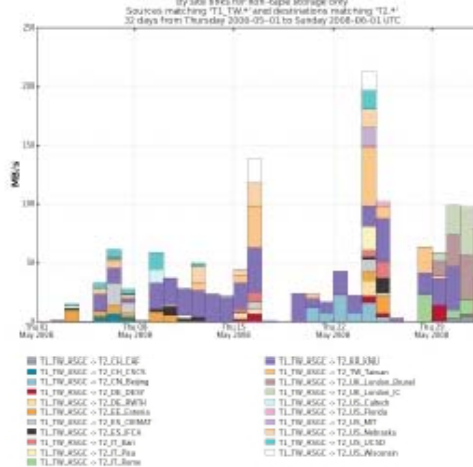
Site name convention for CMS is Tx_[REGION]_[SITENAME] (e.g. "T2_IT_Pisa")

- Countries (e.g. "IT") only include national T2's (e.g. "T2_IT_Pisa")
- Regions (e.g. "CNAF") also include non-national but associated T2's (e.g. "T2_HU_Budapest")

CCRC'08 post-mortem workshop - CERN, 12-13 June 08

T1_TW_ASGC → Sites

Daily CMS PhEDEx transfer rate, Debug + Production



D. Bonacorsi

46



Simon Fraser



Tier-2 Centres in CCRC'08 - CMS

Tier-1 → Tier-2 peak transfer rates

T1→Region best peak	ASGC	CERN	CNAF	FNAL	FZK	IN2P3	PIC	RAL	
ASGC	191.96	-1	102.06	132.51	87.46	52.19	85.21	28.55	
CERN	63.78	240.48	178.93	132.39	187.92	116.92	74.88	62.4	
CNAF	20.98	-1	426.36	75.82	76.21	-1	14.94	17.21	
FNAL	34.84	-1	127.71	763.14	188.73	204.98	75.83	163.22	
FZK	123.72	-1	81.73	119.79	311.6	67.68	6.64	58.28	
IN2P3	42.86	-1	65.71	98.85	97.65	291.7	21.81	-1	
PIC	20.38	-1	48.92	88.51	79.97	44.93	101.54	42.95	
RAL	-1	-1	37.5	38.11	64.71	39.39	57.76	237.71	



Simon Fraser



Tier-2 Centres in CCRC'08 - CMS




Phase-1 : Results




[Phase-1 :: Physics Groups workflows]

[May 5-16]



“controlled” (central) submission

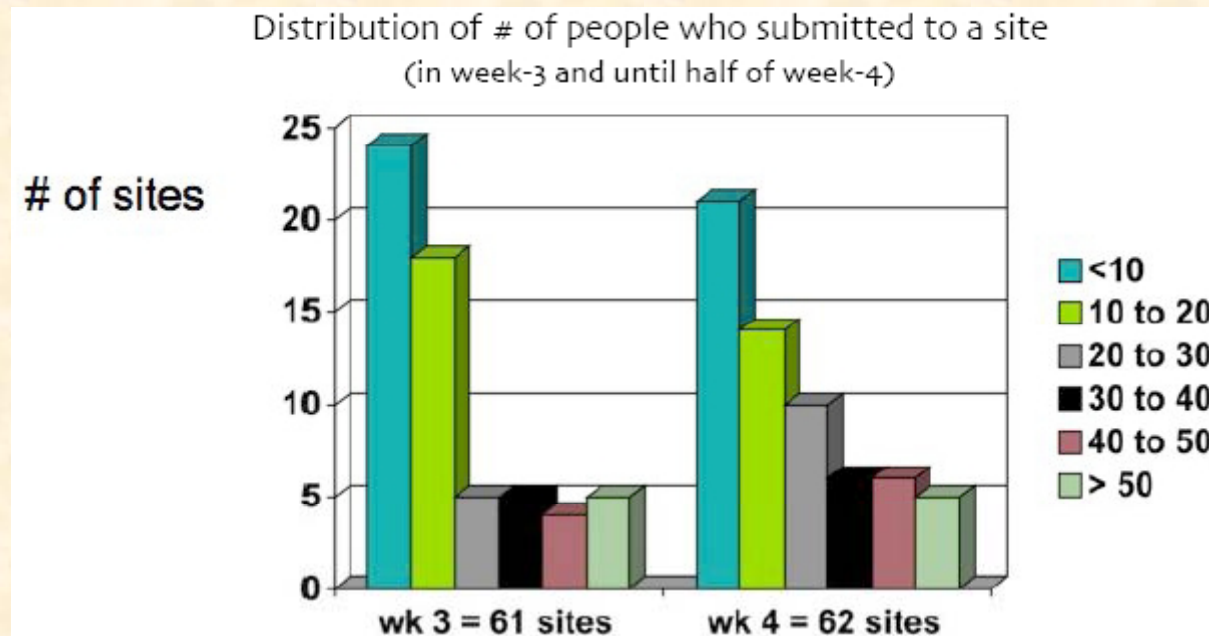


Simon Fraser




Tier-2 Centres in CCRC'08 - CMS

Chaotic analysis mode



Tier-2 Centres in CCRC'08 - CMS

Number of jobs per day



Simon Fraser



Tier-2 Issues/Concerns

- Communications: *Uneven across WLCG. Seems to be good in North America and generally OK within a given country. However, improvements needed for Tier-2s associated with Tier-1s that are not in the same country.*
 - ➔ *work to improve this over the summer: regional Tier-2 coordinators, experiment Tier-2 coordinators, integrate Tier-2s into the GDB as much as practical, improve wiki,...*
- Upcoming onslaught of users: *Some user analysis tests have been done (CMS), but still not on the scale we can expect when data come. Furthermore, new users will be inexperienced.*
- User Support: *Ticketing system exists but it is not really used for user support issues. This affects Tier-2s especially.*

Tier-2 Issues/Concerns

- Better monitoring: Pledges vs actual vs used.
 - set up same system as for Tier-1s. However, this is more difficult: >125 sites vs 11 Tier-1s. Tier-2 coordinators should help with this.
- Hardware acquisitions: Advice on what is best to buy, especially for smaller Tier2s that have less experience.
 - move to SpecInt-2006 ASAP. SI2k is no longer useful; not applicable to new hardware (e.g. large caches).
- Federated Tier-2s: What are the best tools to use to federate sites in to one Tier-2? Priorities, accounting, etc.

How does one account for federated Tier-2 reliability; straight average is often misleading.

Tier-2 Issues/Concerns

- Interoperability of EGEE, OSG, and NDGF should be improved
- Tier-2 capacity: Do we need more resources given the larger size of some of the data sets?
- Software distribution: Could be smoother. Set up Tier-2 software installation coordinators?

Summary

- *The role of the Tier-2 centres has increased markedly in the last year
→ >50% of Monte Carlo simulation is done in the T2s now.*
- *The CCRC'08 exercise is considered a success by the Tier2s and by the experiments.*
- *Availability and reliability are up, but still need improvement.*
- *Resource acquisition vs pledges is better but still needs work*
- *Issues for Tier2s:* - *communication should be improved*
- *work should ramp up on chaotic user analysis*
- *reporting actual resources should be established*
- *improved user support is needed*



presque



Simon Fraser

