




Resource Usage in 2010

R. Graciani
LHCb-Tier1 Jamboree
Lyon 7-8 March 2010

Summary from LHCb-PUB-2011-008

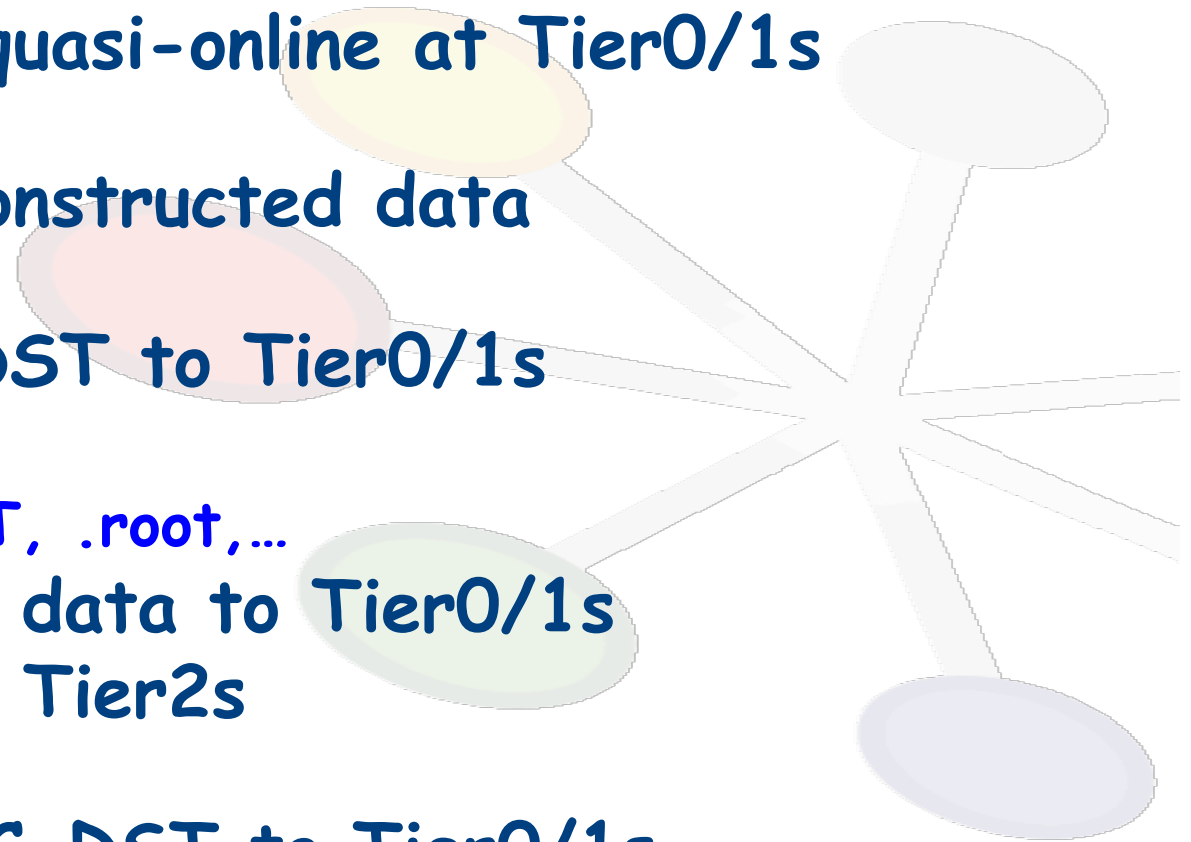


- Pledges for 2010
- LHC conditions
- CPU Usage:
 - WLCG Accounting
 - DIRAC Accounting
- Storage usage
- Summary





- Take data at CERN
 - RAW
- Distribute RAW data to Tier1s
- Reconstruct quasi-online at Tier0/1s
 - SDST
- Stripped reconstructed data
 - DST
- Distributed DST to Tier0/1s
- Analysis
 - mDST, μ DST, .root,...
- Upload USER data to Tier0/1s
- Simulation at Tier2s
 - MC-DST
- Distribute MC-DST to Tier0/1s






2010 Pledges


Date	Site	kHS06	Disk (TB)	Tape (TB)
Apr'10	CERN	23	1290	1800
	Tier-1	43	3254	3036
	Tier-2	42	430	0

Tier1s	CPU		Disk		Tape	
Site	HS06	%	TB	%	TB	%
GRIDKA	7480	17.0%	560	17.2%	408	13.4%
IN2P3	9742	22.1%	728	22.4%	531	17.5%
CNAF	5500	12.5%	450	13.8%	450	14.8%
NIKHEF/ SARA	8992	20.4%	707	21.7%	1012	33.3%
PIC	2632	6.0%	197	6.1%	189	6.2%
RAL	8184	18.6%	612	18.8%	446	14.7%



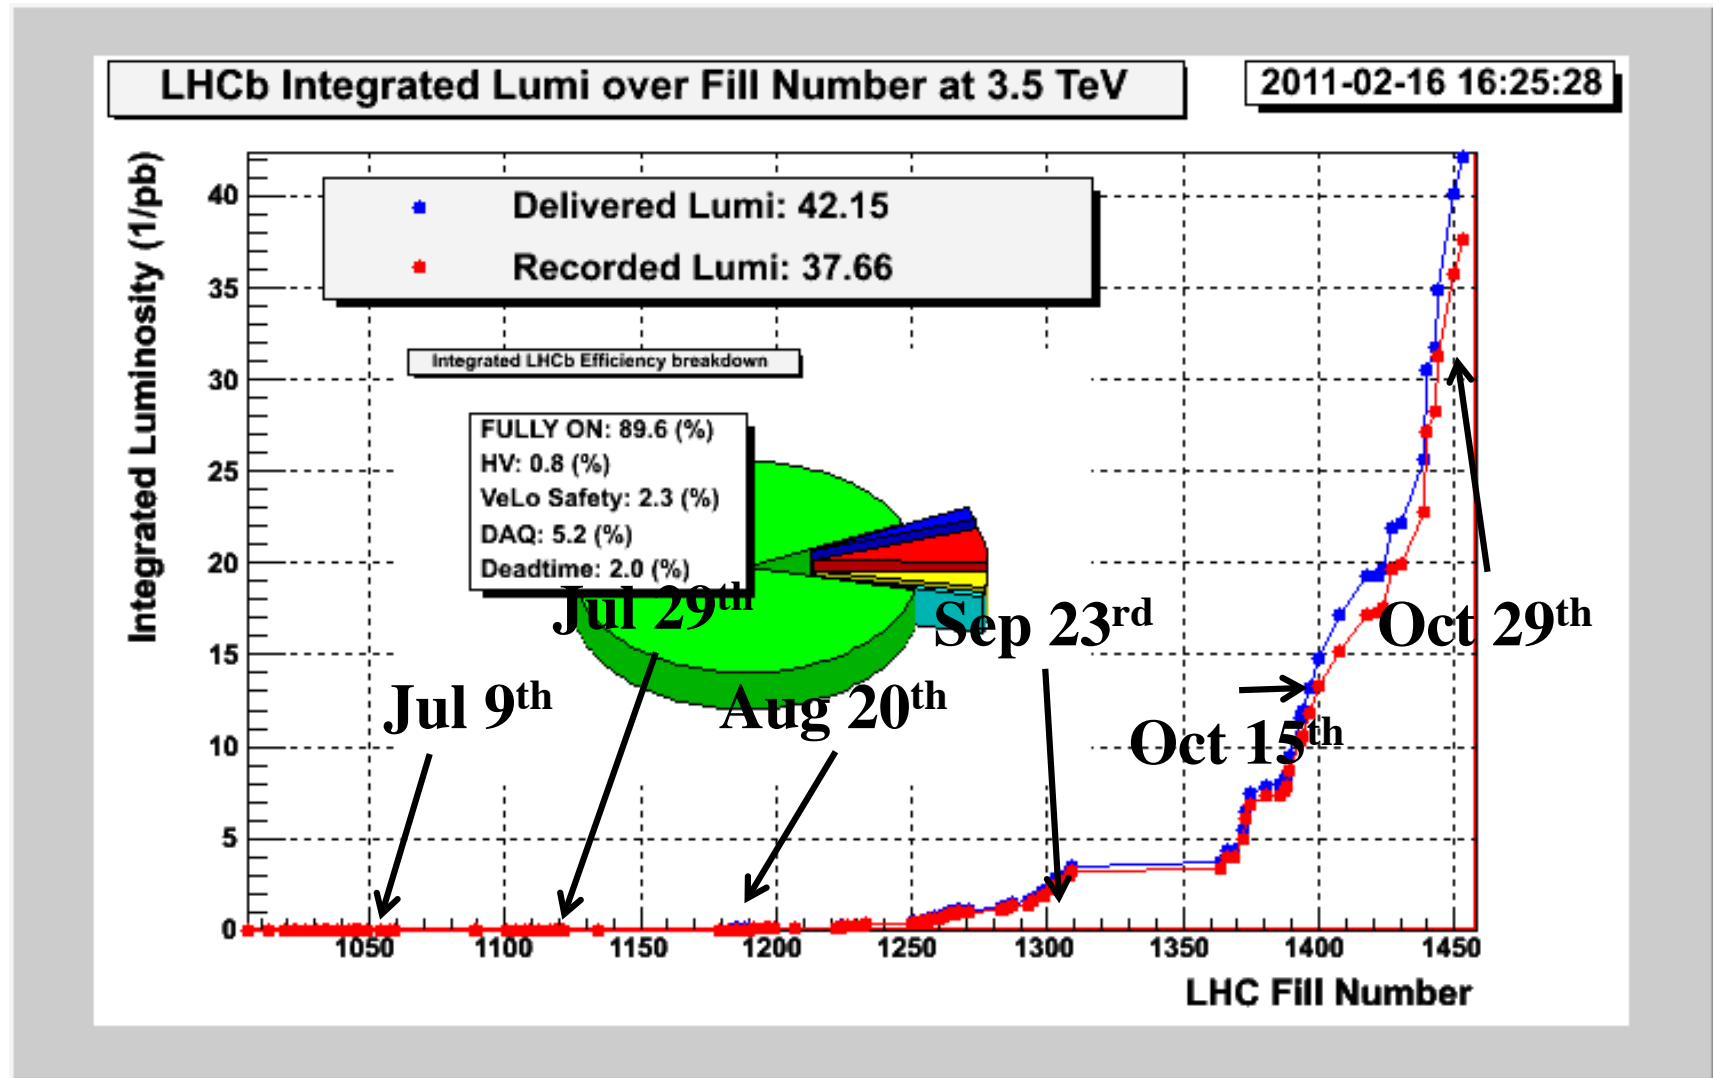


LHCb Nominal: $200 \cdot 10^{30}/\text{cm}^2 \cdot \text{s}$



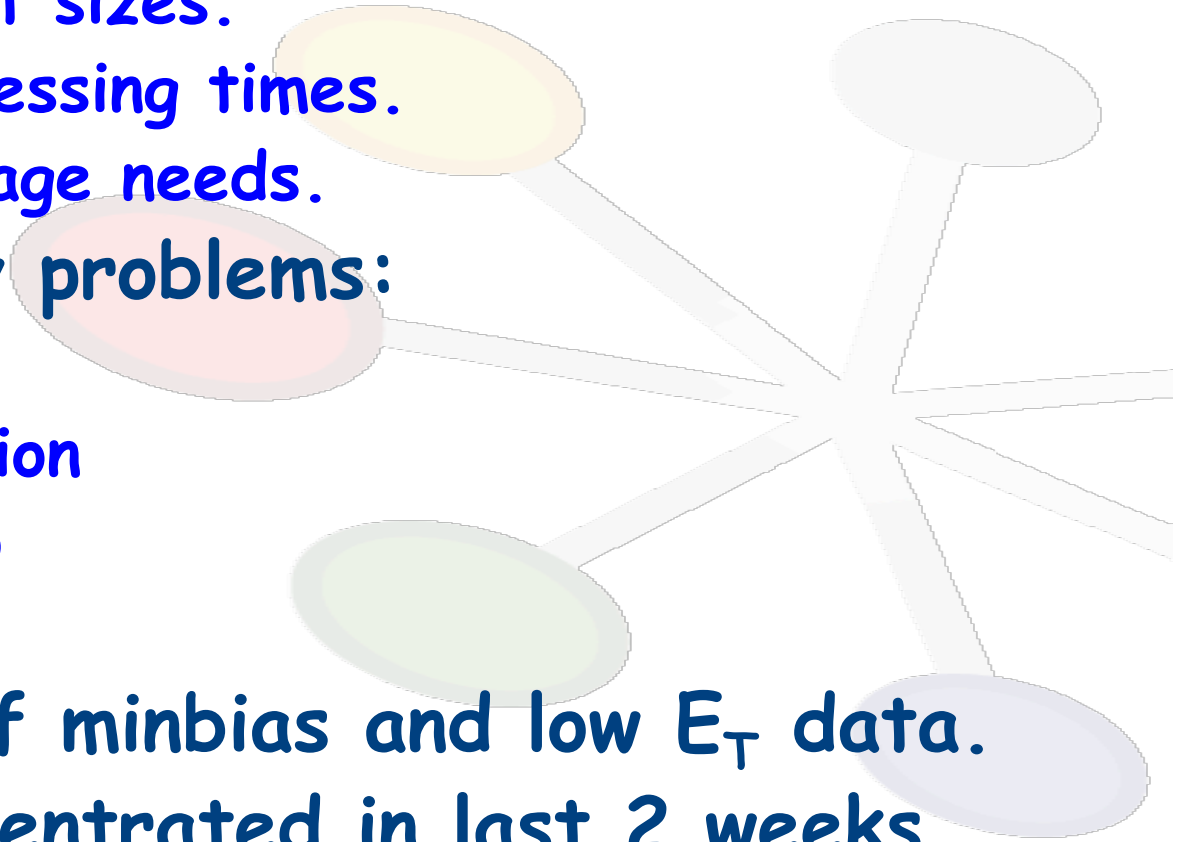


LHCb Nominal: 2000 / pb





- DAQ run all the time at “nominal” rate.
- LHC squeezed bunches to increase lumi:
 - Larger event sizes.
 - Larger processing times.
 - Larger storage needs.
- Caused many problems:
 - Trigger
 - Reconstruction
 - Monte Carlo
 - Stripping
- Took a lot of minbias and low E_T data.
- All lumi concentrated in last 2 weeks.





Collected RAW data

SE	Size (TB)	# of Files
CERN-RAW (T1D0)	97.4	87233
CERN-RDST (T1D1)	83.4	76711
CERN	180.8	163944
CNAF-RAW (T1D0)	19.5	17847
GRIDKA-RAW (T1D0)	27.5	25141
IN2P3-RAW (T1D0)	34.4	31666
NIKHEF-RAW (T1D0)	33.8	31024
PIC-RAW (T1D0)	9.1	8520
RAL-RAW (T1D0)	30.6	27978
Tier1s	154.9	142176

✓ File size ~ 1 GB

✗ 250 TB of RAW were expected for 1 nominal year.




CPU usage from WLCG Accounting

HS06	Norm CPU	Fraction	Norm Elapse	CPU Eff
CERN-Tier0	7166	14.5%	8401	85%
IT-Tier1	3817	21.0%	4598	83%
DE-Tier1	4039	22.3%	4725	85%
FR-Tier1	2437	13.4%	2806	87%
NL-Tier1	3188	17.6%	3775	84%
ES-Tier1	1412	7.8%	1652	86%
UK-Tier1	3255	17.9%	3853	84%
Tier1s	18148	36.6%	21409	85%
Tier2s	24273	48.9%	26469	92%
All	49587	100.0%	56279	88%

- Reasonable contributions from all Sites
- ✓ Very high efficiency



CPU usage DIRAC Accounting





CPU usage DIRAC Accounting

	Reco.	Strip.	MC	Merge	Sam	User
CPU	232,930	1,318	1,071,925	2,143	1,740	323,644
Time	257,744	2,836	1,192,443	31,125	6,218	468,685
Efficiency	90.4%	46.5%	89.9%	6.9%	28.0%	69.1%
Fraction	13.2%	0.1%	60.9%	1.6%	0.3%	23.9%

- ✓ High CPU efficiency.
- ✓ Fast full reprocessing.
- ✗ Low average usage (40% of pledge).
- ✗ MC10 available very late.



Normalization using MC y Reco.

	LHCb (MC)		LHCb* (Reco)		Comparison	Pledge
	HS06	Site/Tier0	HS06·c	Site/Tier0	(L*-L)/L	Site/Tier0
Tier0	6640	100.0%	9019	100.0%	35.83%	100.0%
DE-T1	3005	45.3%	3905	43.3%	29.95%	32.5%
ES-T1	1070	16.1%	1529	17.0%	42.90%	11.5%
FR-T1	1834	27.6%	2602	28.9%	41.88%	42.5%
IT-T1	2440	36.7%	3678	40.8%	50.74%	23.9%
NL-T1	1981	29.8%	2791	30.9%	40.89%	39.1%
UK-T1	2299	34.6%	3058	33.9%	33.01%	35.6%
Tier1s	12745	191.9%	17402	192.9%	36.54%	187.0%

- ✓ Use MC or Reco. to normalize.
- ✓ Very good agreement (< 10%).
- ✓ Reco needs 75% scaling.
- ✓ Same procedure applied for all sites.



Comparison with WLCG

	LHCb		WLCG		Comparison
	HS06	Site/Tier0	HS06	Site/Tier0	(W-L)/L
Tier0	6640	100.0%	7166	100.0%	7.9%
DE-T1	3005	45.3%	4039	56.4%	34.4%
ES-T1	1070	16.1%	1412	19.7%	32.0%
FR-T1	1834	27.6%	2437	34.0%	32.9%
IT-T1	2440	36.7%	3817	53.3%	56.4%
NL-T1	1981	29.8%	3188	44.5%	60.9%
UK-T1	2299	34.6%	3255	45.4%	41.6%
Tier1s	12745	191.9%	18148	253.3%	42.4%
Tier2s	15411	232.1%	24273	338.7%	57.5%


X Large differences in normalization.



Reconstruction Job Failure rates

Reconstruction Job Rates

43 Weeks from Week 13 of 2010 to Week 05 of 2011



Execution Complete	71.3%
Pending Requests	10.7%
Application Finished With Errors	9.0%
Stalled	4.3%
Received Kill signal	1.8%
Watchdog identified this job as stalled	1.6%
Chosen site is not eligible	0.4%
Job has exceeded maximum wall clock time	0.3%
... plus 15 more	


Generated on 2011-02-23 08:48:47 UTC



Simulation Job Failure rates

Simulation Job Rates

43 Weeks from Week 13 of 2010 to Week 05 of 2011

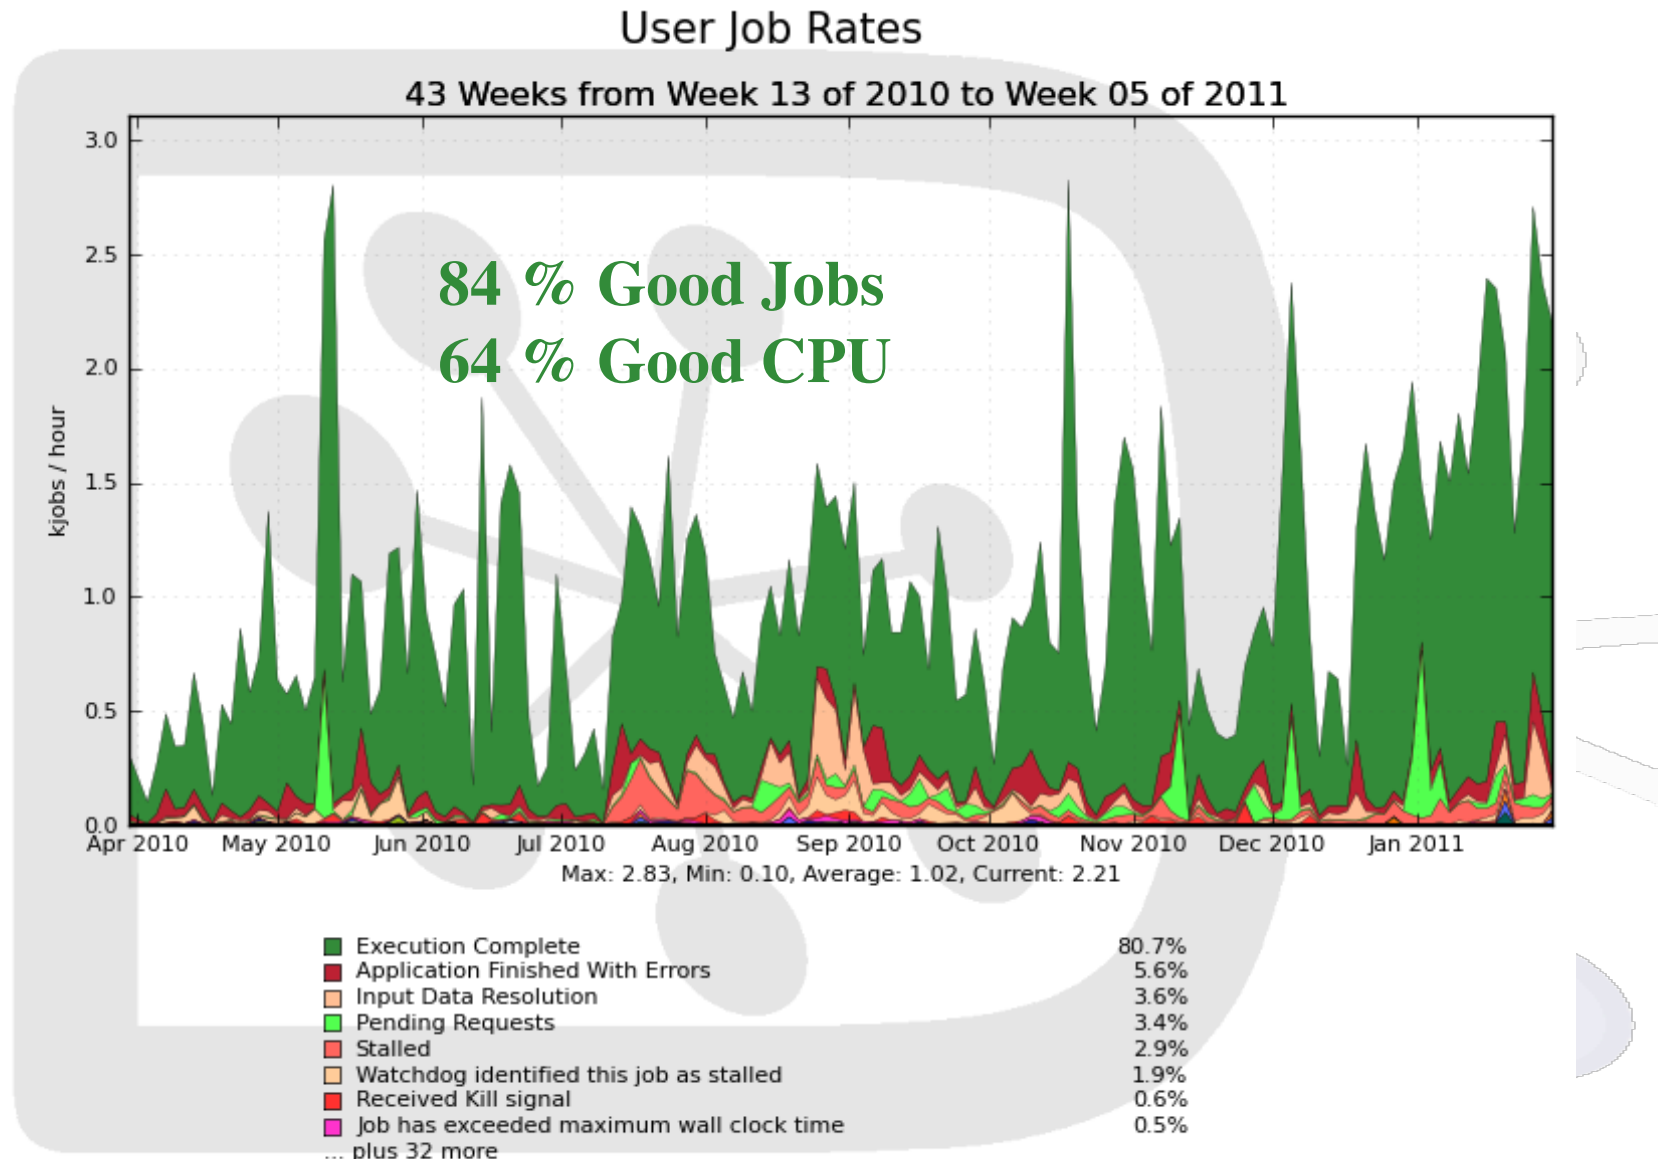


Execution Complete	86.6%	Input Sandbox Download	0.1%
Pending Requests	5.3%	Cannot assign sandbox to job	0.1%
Stalled	4.1%	Job has reached the CPU limit of the queue	0.0%
Application Finished With Errors	2.3%	Exception During Execution	0.0%
Received Kill signal	0.8%	Job has insufficient disk space to continue	0.0%
Watchdog identified this job as stalled	0.3%	Cannot insert in task queue	0.0%
Job has exceeded maximum wall clock time	0.2%	Job Wrapper Initialization	0.0%
Uploading Job Outputs	0.2%	Job exceeded maximum load average	0.0%

Generated on 2011-02-23 08:47:34 UTC

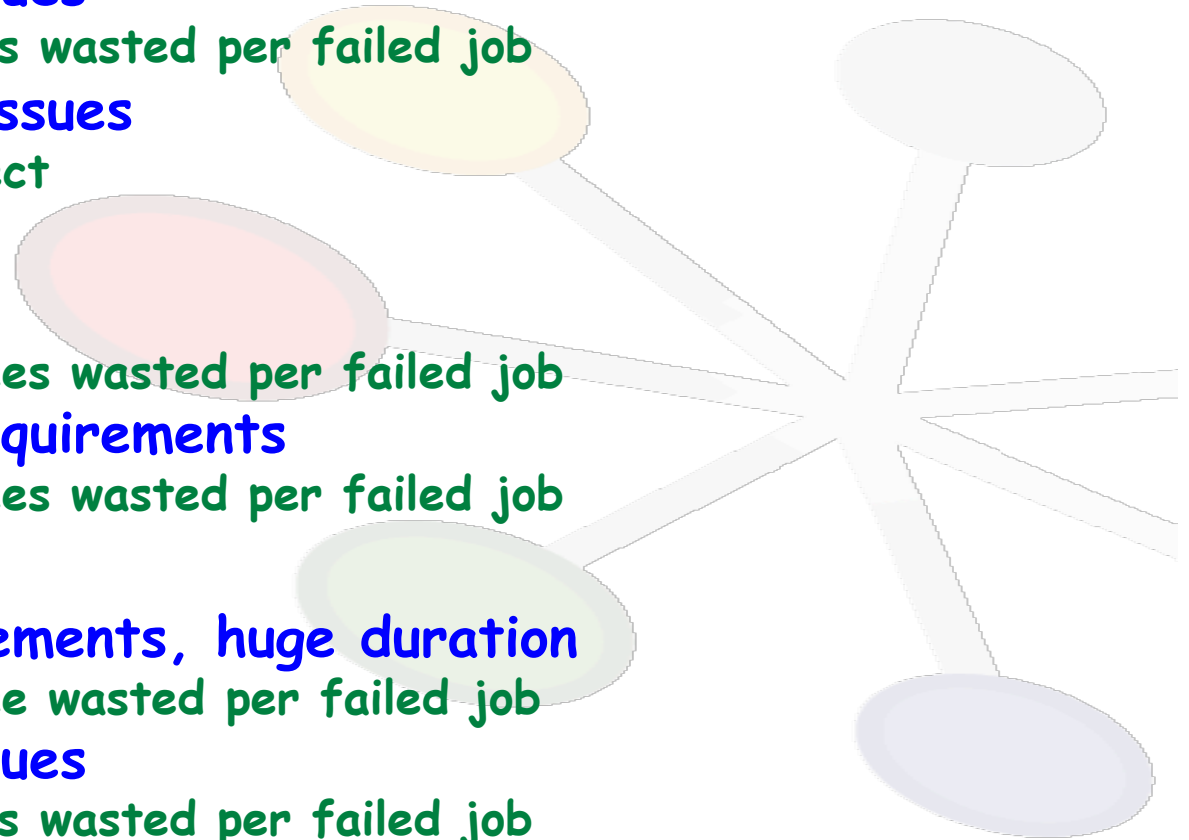


Reconstruction Job Failure rates





- **Reconstruction:**
 - **Too long execution time, too large memory**
 - ☆ Many resources wasted per failed job
 - **Input data issues**
 - ☆ Few resources wasted per failed job
 - **Shared area issues**
 - ☆ Variable effect
- **Simulation:**
 - **Infinite loops**
 - ☆ Many resources wasted per failed job
 - **Wrong CPU requirements**
 - ☆ Many resources wasted per failed job
- **Users:**
 - **Wrong Requirements, huge duration**
 - ☆ Many resource wasted per failed job
 - **Input data issues**
 - ☆ Few resources wasted per failed job





Disk Summary	Pledge (TB)	Seen by SLS		Seen by LHCb	
20/12/2010		TB		TB	
		Total	Used	Used	Pledge-Used
FZK	495	500	331	339.9	155.1
IN2P3	610	641	334	320.7	289.3
CNAF	450	463	392	391.6	58.4
NL-T1	560	563	339	254.5	305.5
PIC	240	255	138	138.3	101.7
RAL	505	791	562	453.3	51.7
Tier1s	2860	3213	2096	1897.5	962.5
CERN	1135	1175	922	763.6	371.4

✓ There is space available.

✗ It is highly fragmented, many full SEs.





Tape Usage

Tape Summary	Pledge (TB)	Seen by SLS		Seen by LHCb	
20/12/2010		TB		TB	
		Total	Used	Used	Pledge-Used
FZK	350			160.7	189.3
IN2P3	555			188.4	366.6
CNAF	265			126.4	138.6
NL-T1	420			161.1	258.9
PIC	130			65	65
RAL	380			201.6	178.4
Tier1s	2100			903.2	1196.8
CERN	1635			844.7	790.3



- LHC has setup has caused problems:
 - Issues with trigger, solved
 - Issues with reconstruction, solved
 - Issues with stripping, solved
 - Issues with data volumes, partially solved
- Not full use of pledge resources, but
 - Good rate of successful jobs
 - Good efficiency in usage
 - Can achieve large peaks of CPU
- Issues with normalization
- Need to work on
 - Software access
 - Data access
 - Data management
- Hope to improve collaboration with YOU







Full CPU Norm. Comparison

	LHCb		WLCG		LHCb*		Comparison	
	HS06	Site/ Tier0	HS06	Site/ Tier0	HS06.c	Site/ Tier0	(W-L)/L	(L*-L)/L
Tier0	6640	100.0%	7166	100.0%	9019	100.0%	7.9%	35.83%
DE-T1	3005	45.3%	4039	56.4%	3905	43.3%	34.4%	29.95%
ES-T1	1070	16.1%	1412	19.7%	1529	17.0%	32.0%	42.90%
FR-T1	1834	27.6%	2437	34.0%	2602	28.9%	32.9%	41.88%
IT-T1	2440	36.7%	3817	53.3%	3678	40.8%	56.4%	50.74%
NL-T1	1981	29.8%	3188	44.5%	2791	30.9%	60.9%	40.89%
UK-T1	2299	34.6%	3255	45.4%	3058	33.9%	41.6%	33.01%
Tier1s	12745	191.9%	18148	253.3%	17402	192.9%	42.4%	36.54%
CH-T2	441	6.6%	771	10.8%			74.8%	
DE-T2	549	8.3%	2114	29.5%			285.1%	
ES-T2	782	11.8%	830	11.6%			6.1%	
FR-T2	2164	32.6%	2553	35.6%			18.0%	
IT-T2	2822	42.5%	5739	80.1%			103.4%	
RO-T2	101	1.5%	238	3.3%			135.6%	
RU-T2	1656	24.9%	2610	36.4%			57.6%	
UK-T2	6514	98.1%	9420	131.5%			44.6%	
Tier2s	15411	232.1%	24273	338.7%			57.5%	





Usage of the CPU pledge

HS06	Pledge	WLCG	LHCb	WLCG(%)	LHCb(%)
CERN	23000	7166	6640	31%	29%
CNAF	5500	3817	2440	69%	44%
GRIDKA	7480	4039	3005	53%	40%
IN2P3	9742	2437	1834	25%	19%
NL-T1	8992	3188	1981	35%	22%
PIC	2632	1412	1070	54%	41%
RAL	8184	3255	2299	40%	28%

- **Average work April 2010-Jan 2010**