LHCb Resource Needs

LHCC review – Jan'05

Ricardo Graciani

Overview

- The LHCb Computing Model
  - CERN, Tier1's & Tier2's
- 2008 Requirements
  - RAW data
  - Reconstruction
  - Stripping and Reprocessing
  - Monte Carlo Simulation
  - CPU
  - Storage
- 2009 Requirements
- 2010 Requirements
- Summary

LHCb Computing Model

	Efficiency
Production CPU	85%
Analysis CPU	60%
Disk Storage	70%
MSS Storage	100%

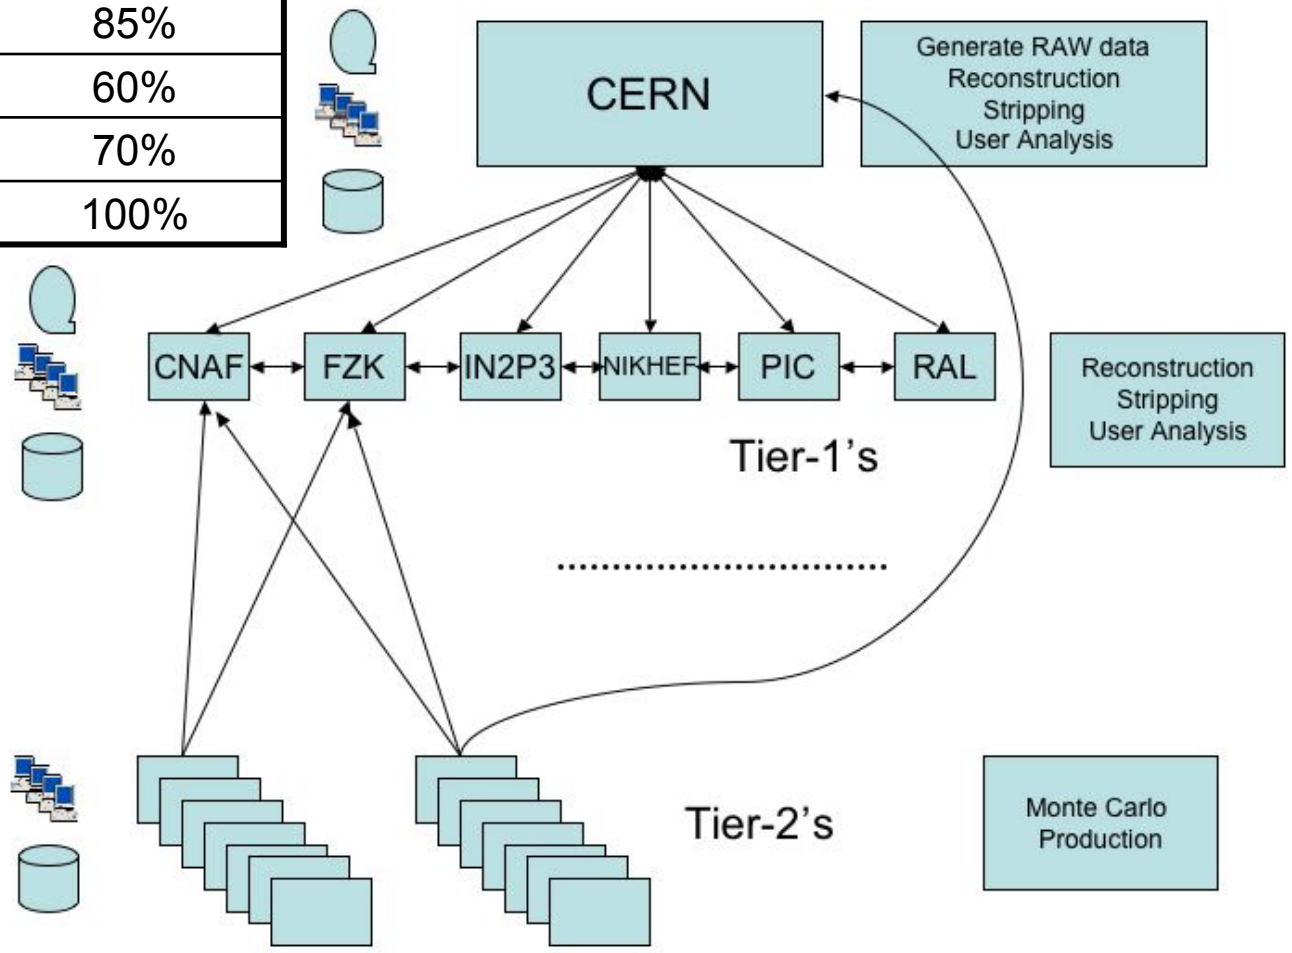

CERN (Tier0+Tier1)

- Get “Master copy” of RAW data.
- Contributes to:
  - Reconstruction
  - Stripping
  - User Analysis
- Keeps complete RAW, DST and MC data samples on MSS.

Tier1's

- 6 Candidates:
  - CNAF (Bologna, Italy)
  - FZK (Karlsruhe, Germany)
  - IN2P3 (Lyon, France)
  - NIKHEF (Amsterdam, Netherlands)
  - PIC (Barcelona, Spain)
  - RAL (UK)
- Contribute to:
  - Reconstruction
  - Stripping
  - User Analysis
- Shared a copy of RAW, DST and MC samples on MSS.
- Master copy of rDST data shared over CERN and Tier1's MSS.
- Different capacities are foreseen.

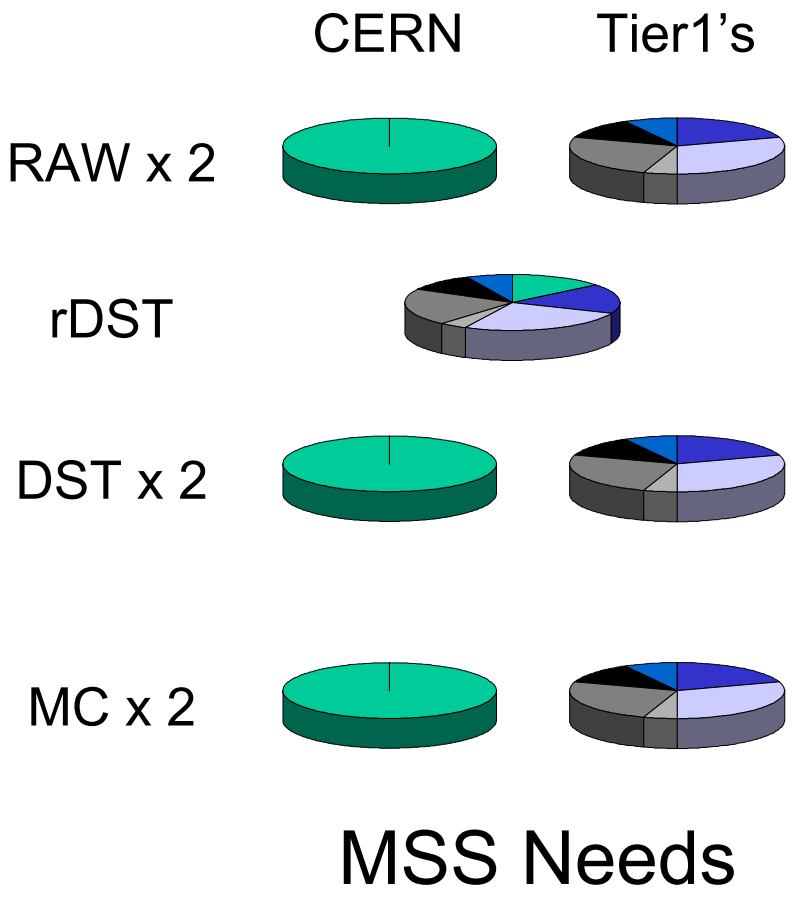

MSS Needs

Tier2's

- 14 Candidates:
- Responsibility of Monte Carlo simulation.
- Transfer output DST to CERN and Tier1's.
- Larger sites may contribute to User Analysis.

2008

- Assumed first year of full data taking:
  - 10^7 seconds @ 2 x $10^{32}cm^{-2}s^{-1}$.
  - Extended over 7 month (April-October)
  - These are “stable running conditions”.
- Data sample:

	b-exclusive	dimuon	D*	b-inclusive	Total
Trigger Rate (Hz)	200	600	300	900	2000
Events (10^9)	2	6	3	9	20

RAW Data

- Produced @ Online FARM.
  - 50 MB/s
- Master copy sent to CERN MSS ("permanent" storage).
  - 500 TB
- 2nd copy distributed over Tier1's MSS.
  - 27 MB/s
  - 500 TB
- Input data for Reconstruction.

Online Reconstruction

- Online Farm
  - 1800 CPUs @ 3 kSI2k (PASTA 2006-07) => 5.4 MSI2k
  - 40 TB of Disk
  - HTL + Online Reconstruction
- 10% of RAW data (b-exclusive):
  - Reconstructed @ Online FARM.
  - 10 MB/s
  - Produced data kept on disk (“fast-access” storage) @ CERN for Monitoring and Calibration purposes.

Reconstruction

- 1st pass during Data Taking @ CERN + Tier1's (7 months).
  - Data Rate: Input 27 MB/s
    Output 27 MB/s
  - CPU Power: CERN 0.44 MSI2k
    Tier1's 2.64 MSI2k
- 2nd pass during winter shut-down @ CERN + Tier1's + Online FARM (2 months).
  - Data Rate: Input: 95 MB/s
    Output: 95 MB/s
  - CPU Power: CERN 0.90 MSI2k
    Tier1's 5.40 MSI2k
- Output to MSS distributed over CERN & Tier1's
  - 500 TB / pass

Stripping

- 4 passes:
  - during Data Taking @ CERN + Tier1's (7 months).
    - Data Rate: Input: 54 MB/s; Output: 8 MB/s
    - CPU Power: CERN 0.08 MSI2k; Tier1's 0.50 MSI2k
  - after Data Taking @ CERN + Tier1's (1 month).
    - Data Rate Input: 381 MB/s; Output: 53 MB/s
    - CPU Power: CERN 0.59 MSI2k; Tier1's 3.51 MSI2k
  - during shut-down @ CERN + Tier1's + Online FARM (2 months).
    - Data Rate Input: 190 MB/s; Output: 26 MB/s
    - CPU Power: CERN 0.17 MSI2k; Tier1's 1.01 MSI2k
  - before next year Data Taking @ CERN + Tier1's (1 month).
- Output distributed to CERN + Tier1's (+ Tier2's) for User Analysis:
  - 140 TB / pass / copy.
  - MSS: 1 copy @ CERN + 1 copy distributed over Tier1's.
  - Disk: copy @ CERN + all Tier1's (+Tier2's) . Older versions removed (2 versions kept for b samples).

Monte Carlo Simulation

- Produced @Tier2's
  - Data Rate: Output 5 MB/s
  - CPU Power: Tier2's 7.65 MSi2k
- Output distributed to CERN + Tier1's (+Tier2's) for User Analysis:
  - 160 TB / copy
  - MSS: 1 copy @ CERN + 1 copy distributed over Tier1's.
  - Disk: current year sample 1 copy @ CERN + 3 copies Tier1's (+Tier2's).

CPU Requirements

Breakdown of CPU needs
(with efficiencies)

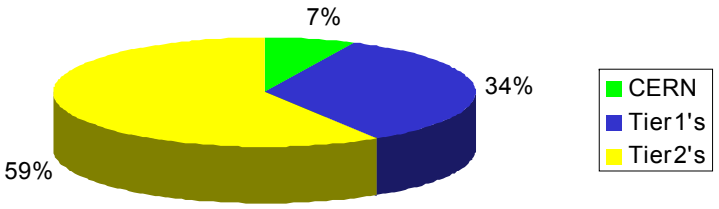

2008

Breakdown of CPU needs
(with efficiencies)

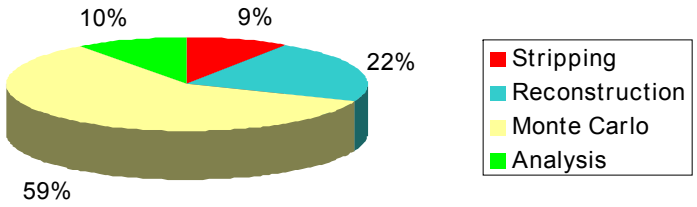

* Online FARM resources not presented here

MS2k years	CERN	Tier1's	Tier2's	Total
Stripping	0.17	1.03	0.00	1.20
Recons.	0.40	2.42	0.00	2.83
Monte Carlo	0.00	0.00	7.65	7.65
Analysis	0.32	0.97	0.00	1.29
Total	0.90	4.42	7.65	12.97

CPU Efficiencies:

Production CPU: 85 % Analysis CPU: 60 %

Permanent Storage (MSS)

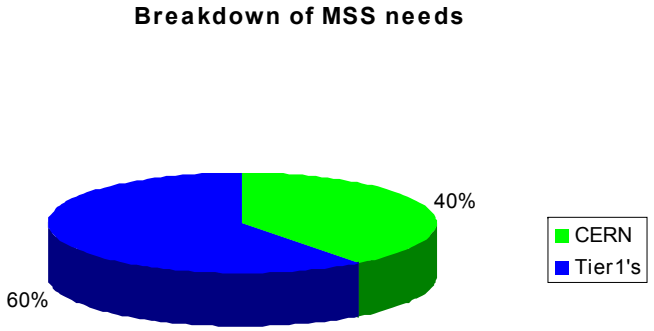

2008

TB	CERN	Tier1's	Total
RAW	500	500	1000
rDST	143	857	1000
Data DST	556	556	1112
MC DST	160	160	321
Total	1359	2074	3433

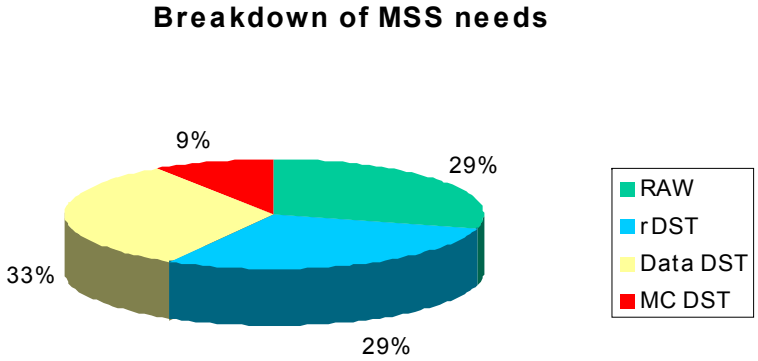

Fast Storage (Disk)

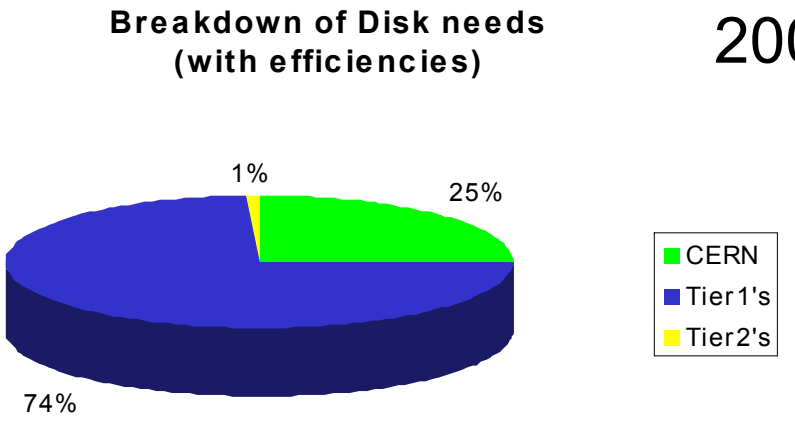

2008

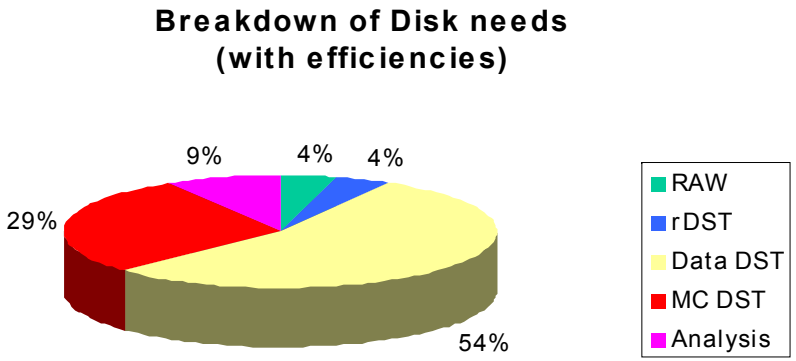

TB	CERN	Tier1's	Tier2's	Total
RAW	136	0	0	136
rDST	136	0	0	136
Data DST	256	1534	0	1790
MC DST	229	687	23	939
Analysis	70	210	0	280
Total	826	2432	23	3281

2009

- Assume constant luminosity by beam de-focussing:
  - 10^7 seconds @ 2 x 10^{32}cm^{-2}s^{-1}.
  - Extended over 7 month (April-October)
- 2008 stripped data is reprocessed during shutdown after 2009 data taking.
- MSS Storage need extends with new data.

2009

- Disk Storage, copies of 2008 data latest version:

	CERN	Tier1's
b-exclusive	1	2
D*	1	1
Di-muon	1	1
b-inclusive	1	2

- Analysis needs increase with data sample.

2010

- Same DAQ conditions as previous years.
- 2008 data fully reconstruct + strip from RAW during data taking.
- 2009 stripped data is reprocessed during shutdown after 2010 data taking.
- MSS Storage need extends with new data.

2010

- Disk Storage, copies of 2008 & 2009 data latest version:

	CERN	Tier1's
b-exclusive	1	2
D*	1	1
Di-muon	1	1
b-inclusive	1	2

- Analysis needs increase with data sample.

Summary (CPU)

MSi2k*year	2006	2007	2008	2009	2010
CERN T0 + T1	0.27	0.54	0.90	1.25	1.88
Tier1s	1.33	2.65	4.42	5.55	8.35
Tier2s	2.29	4.59	7.65	7.65	7.65
Total	3.89	7.78	12.97	14.45	17.87

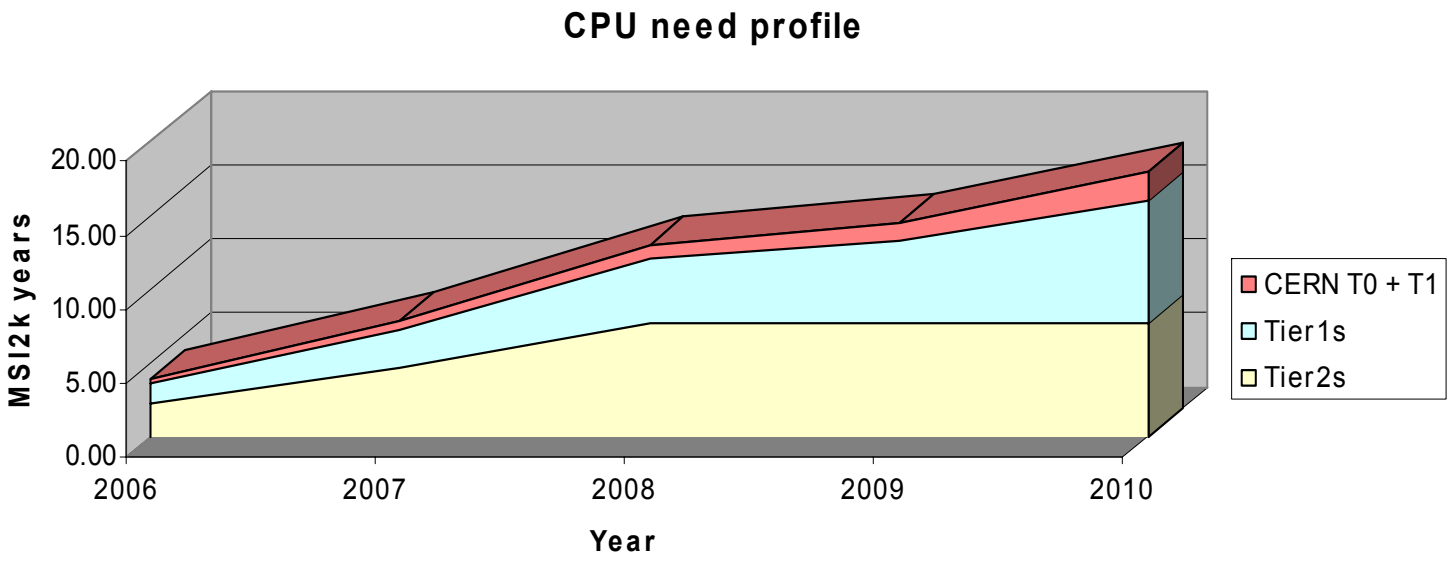

Summary (CPU Power I)

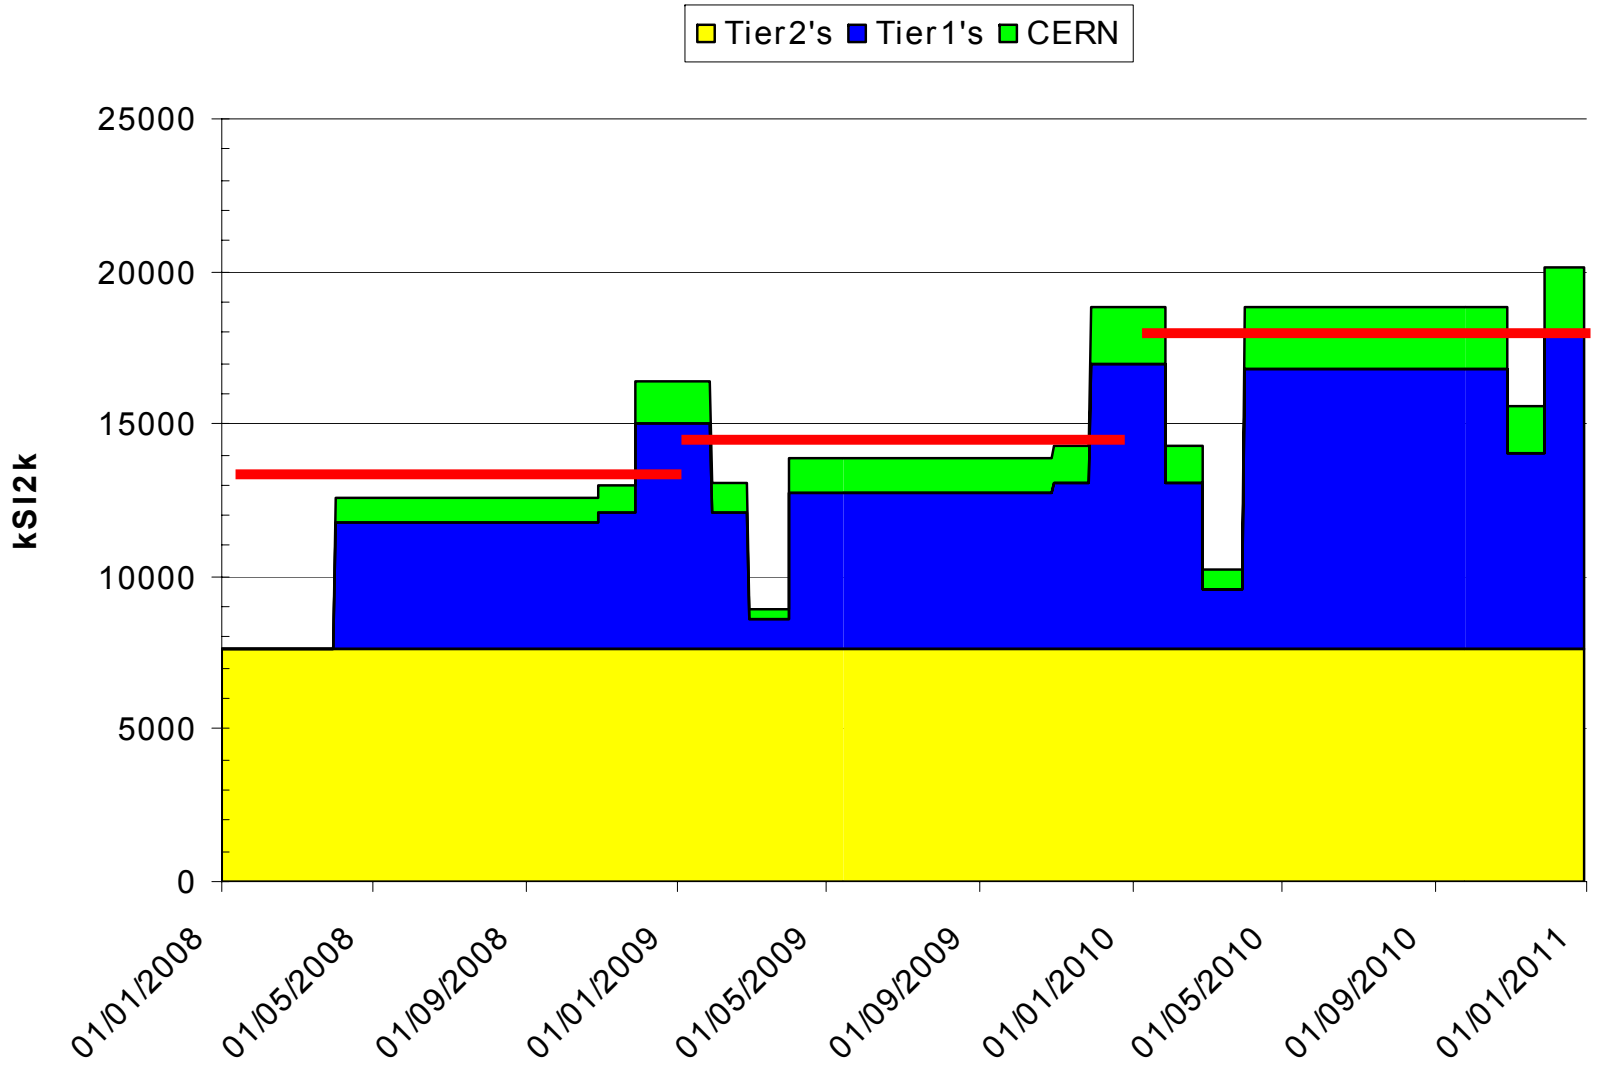

Summary (CPU Power II)

(MSI2k)	Months	Tier2's	Tier1's	CERN
2008	Data taking	7.6	4.1	0.8
	Stripping	7.6	4.5	0.9
	Re-recons	7.6	7.4	1.4
2009	Data taking	7.6	5.1	1.2
	Stripping	7.6	5.4	1.2
	Re-recons	7.6	9.3	1.9
2010	Data taking	7.6	9.2	2.0
	Stripping	7.6	6.4	1.6
	Re-recons	7.6	10.3	2.2

- Analysis assumed to be on going – flat during the year

Summary (MSS)

(TB)	2006	2007	2008	2009	2010
CERN T0 + T1	408	816	1359	2858	4566
Tier1s	622	1244	2074	4286	7066
Tier2s					
Total	1030	2060	3433	7144	11632

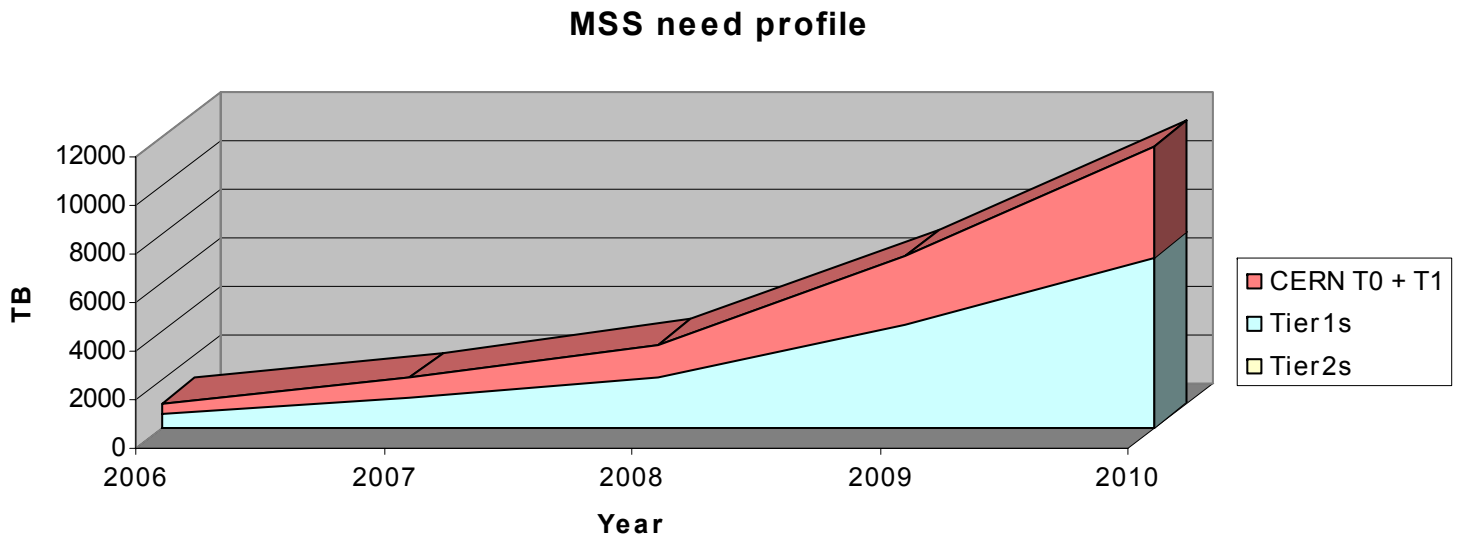

Summary (Disk)

(TB)	2006	2007	2008	2009	2010
CERN T0 + T1	248	496	826	1095	1363
Tier1s	730	1459	2432	2897	3363
Tier2s	7	14	23	23	23
Total	984	1969	3281	4015	4749

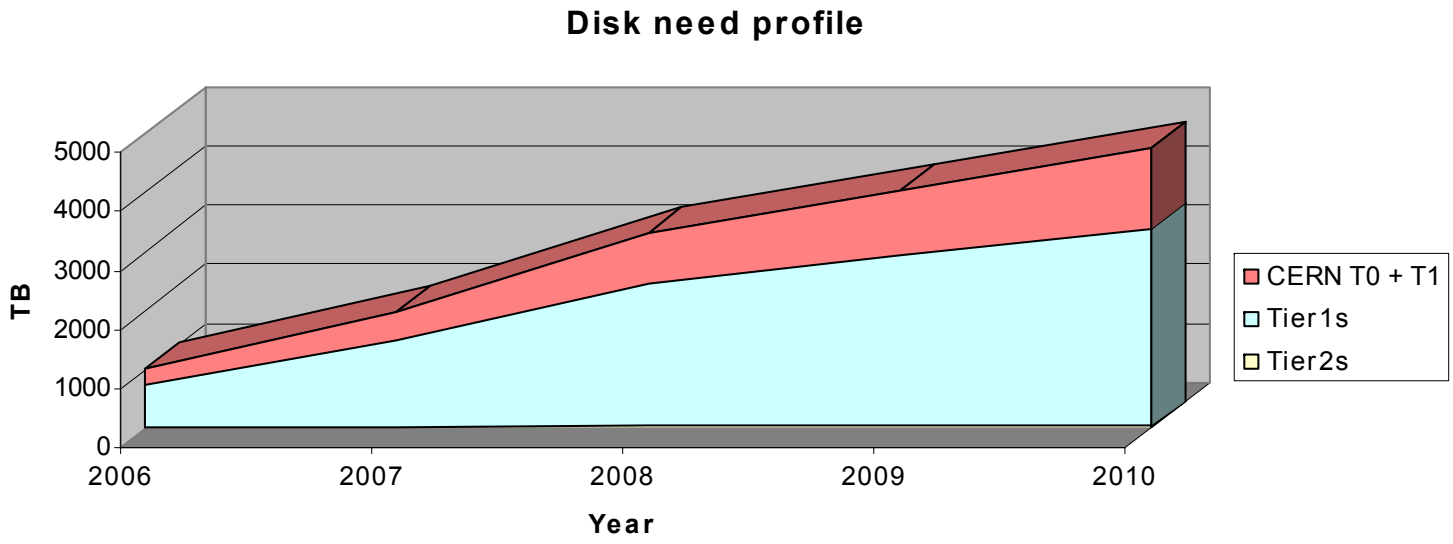

Bandwidth Out

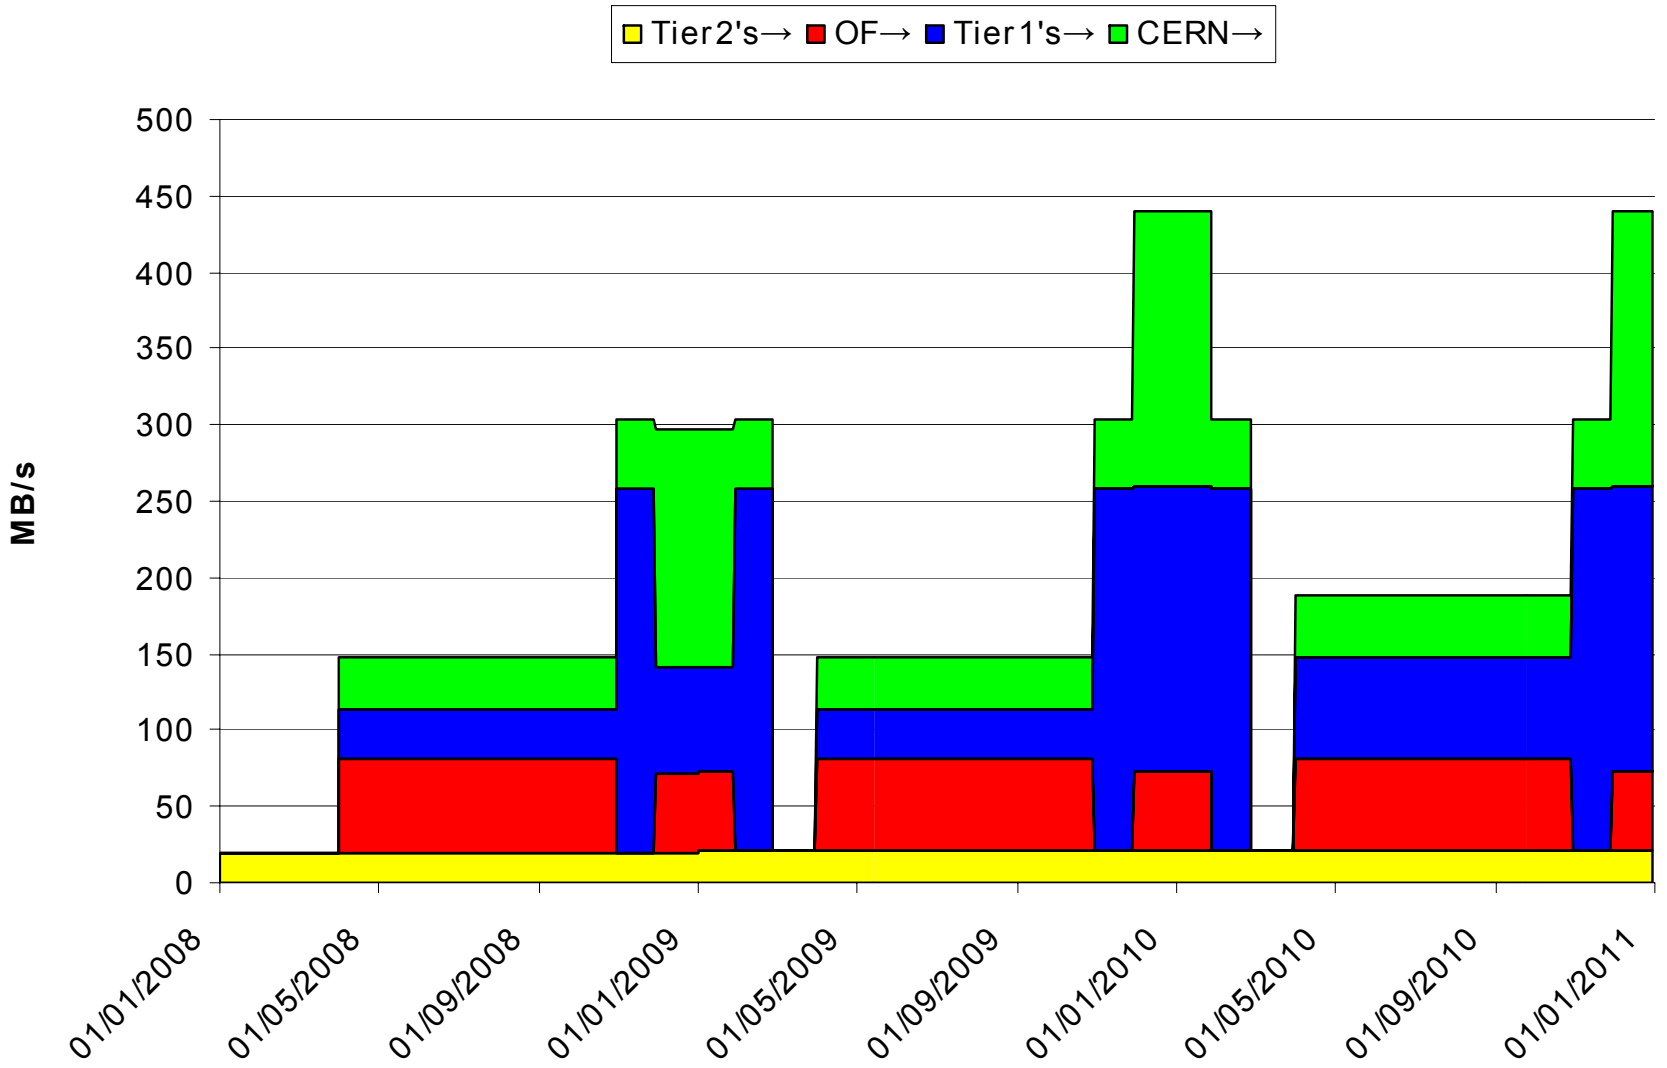

Bandwidth Out

(MB/s)	Months	Tier2's→	Tier1's→	OF→	CERN→
2008	Data taking	20	34	60	34
	Stripping	20	237	0	46
	Reprocessing	20	69	52	157
2009	Data taking	20	34	60	34
	Stripping	20	237	0	46
	Reprocessing	20	187	52	180
2010	Data taking	20	67	60	40
	Stripping	20	237	0	46
	Reprocessing	20	187	52	180

Bandwith In

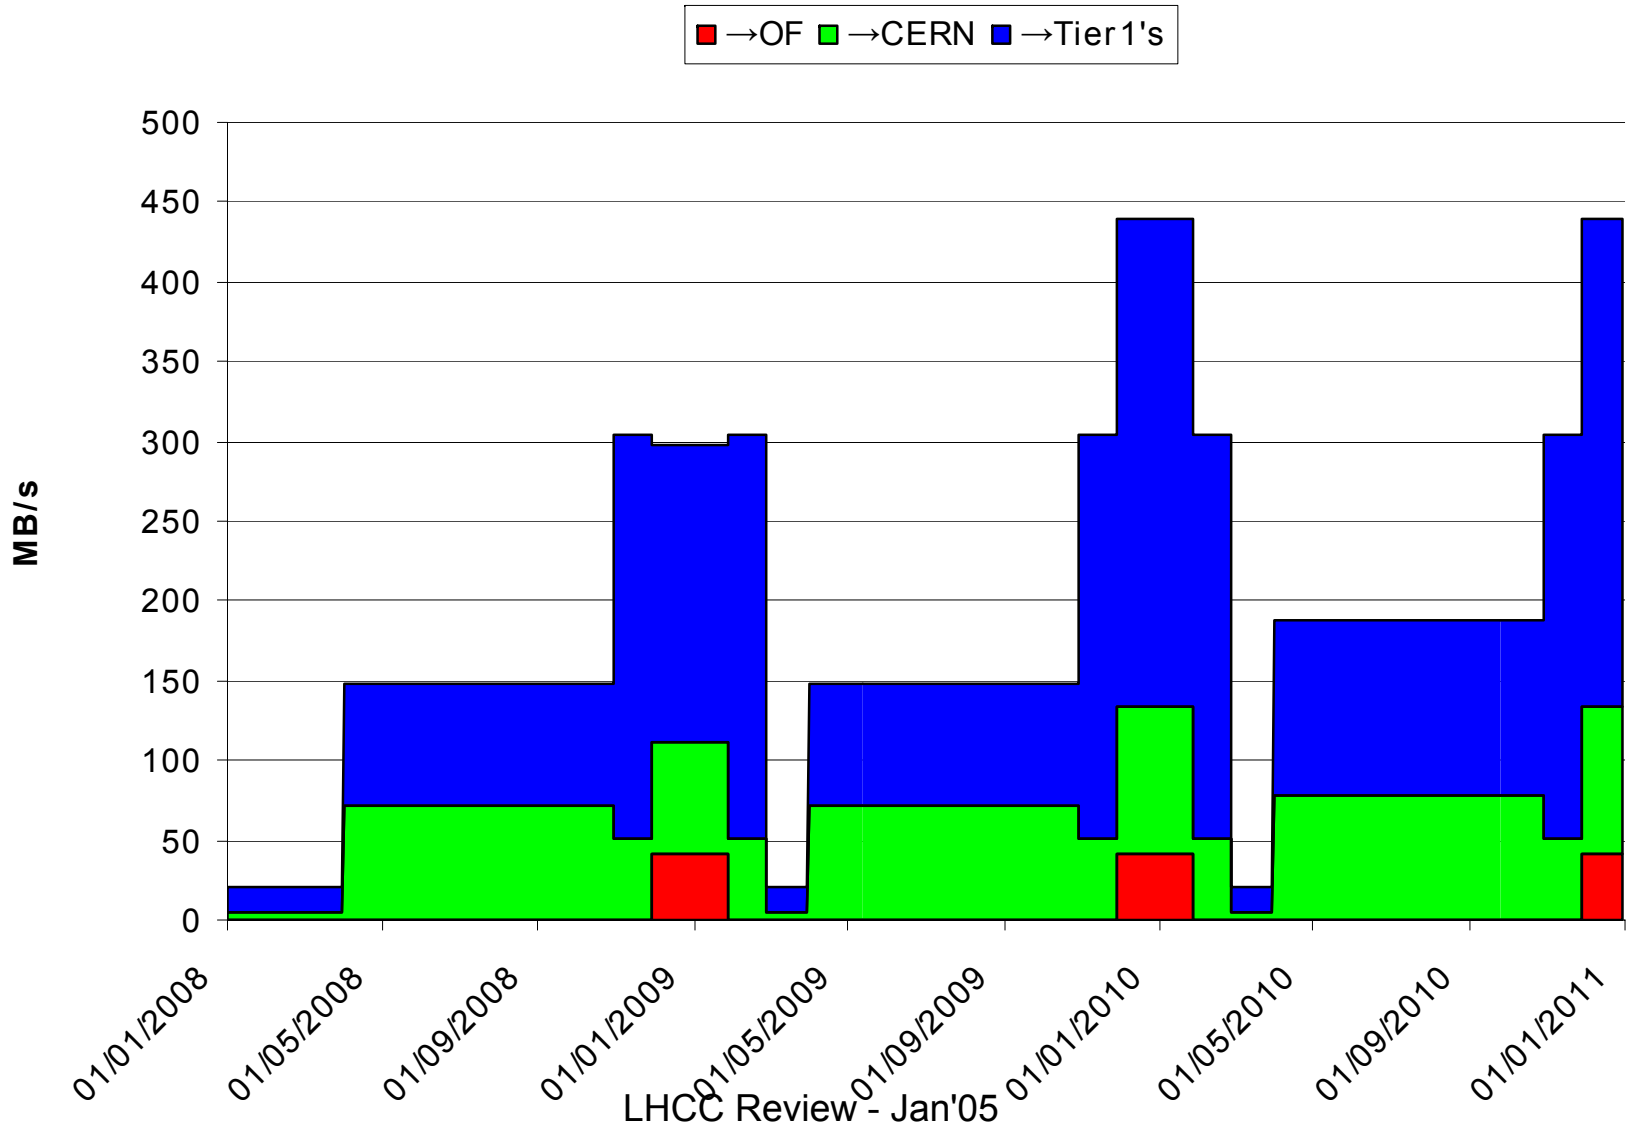
→OF
→CERN
→Tier1's
MB/s
500
450
400
350
300
250
200
150
100
50
0
01/01/2008
01/05/2008
01/09/2008
01/01/2009
01/05/2009
01/09/2009
01/01/2010
01/05/2010
01/09/2010
01/01/2011

Bandwitdh In

(MB/s)	Months	→OF	→Tier1's	→CERN
2008	Data taking	0	76	72
	Stripping	0	253	51
	Reprocessing	41	187	70
2009	Data taking	0	76	72
	Stripping	0	253	51
	Reprocessing	41	305	93
2010	Data taking	0	110	78
	Stripping	0	253	51
	Reprocessing	41	305	93

Bandwidth with FARM

(MB/s)	Months	OF→CERN	CERN→OF
2008	Data taking	60	0
	Stripping	0	0
	Reprocessing	52	41
2009	Data taking	60	0
	Stripping	0	0
	Reprocessing	52	41
2010	Data taking	60	0
	Stripping	0	0
	Reprocessing	52	41

Bandwidth Tier1's

(MB/s)	Months	CERN→Tier1's	Tier1's→CERN	Tier1's→Tier1
2008	Data taking	34	7	5
	Stripping	46	46	38
	Reprocessing	116	13	11
2009	Data taking	34	7	5
	Stripping	46	46	38
	Reprocessing	139	36	30
2010	Data taking	40	13	11
	Stripping	46	46	38
	Reprocessing	139	36	30

Bandwidth Tier2's

(MB/s)	Months	Tier2's→CERN	Tier2's→Tier1's
2008	Data taking	5	15
	Stripping	5	15
	Reprocessing	5	15
2009	Data taking	5	15
	Stripping	5	15
	Reprocessing	5	15
2010	Data taking	5	15
	Stripping	5	15
	Reprocessing	5	15

Summary

	2006	2007	2008	2009	2010
CPU(MSI2k year)	3.89	7.78	12.97	14.45	17.87
MSS (TB)	1030	2060	3433	7144	11632
Disk (TB)	984	1969	3281	4015	4749

- Peak during short stripping and reconstruction: 300, 450, 450 MB/s (total).
- < 100 MB/s Farm $\leftrightarrow$ CERN.
- < 200 MB/s out of CERN (including FARM).
- < 100 MB/s into CERN (including FARM).
- < 250 MB/s out of all Tier1.
- < 60 MB/s into any Tier1.