


WLCG Status

Ian Bird; WLCG Project Leader

LHCC Closed Session

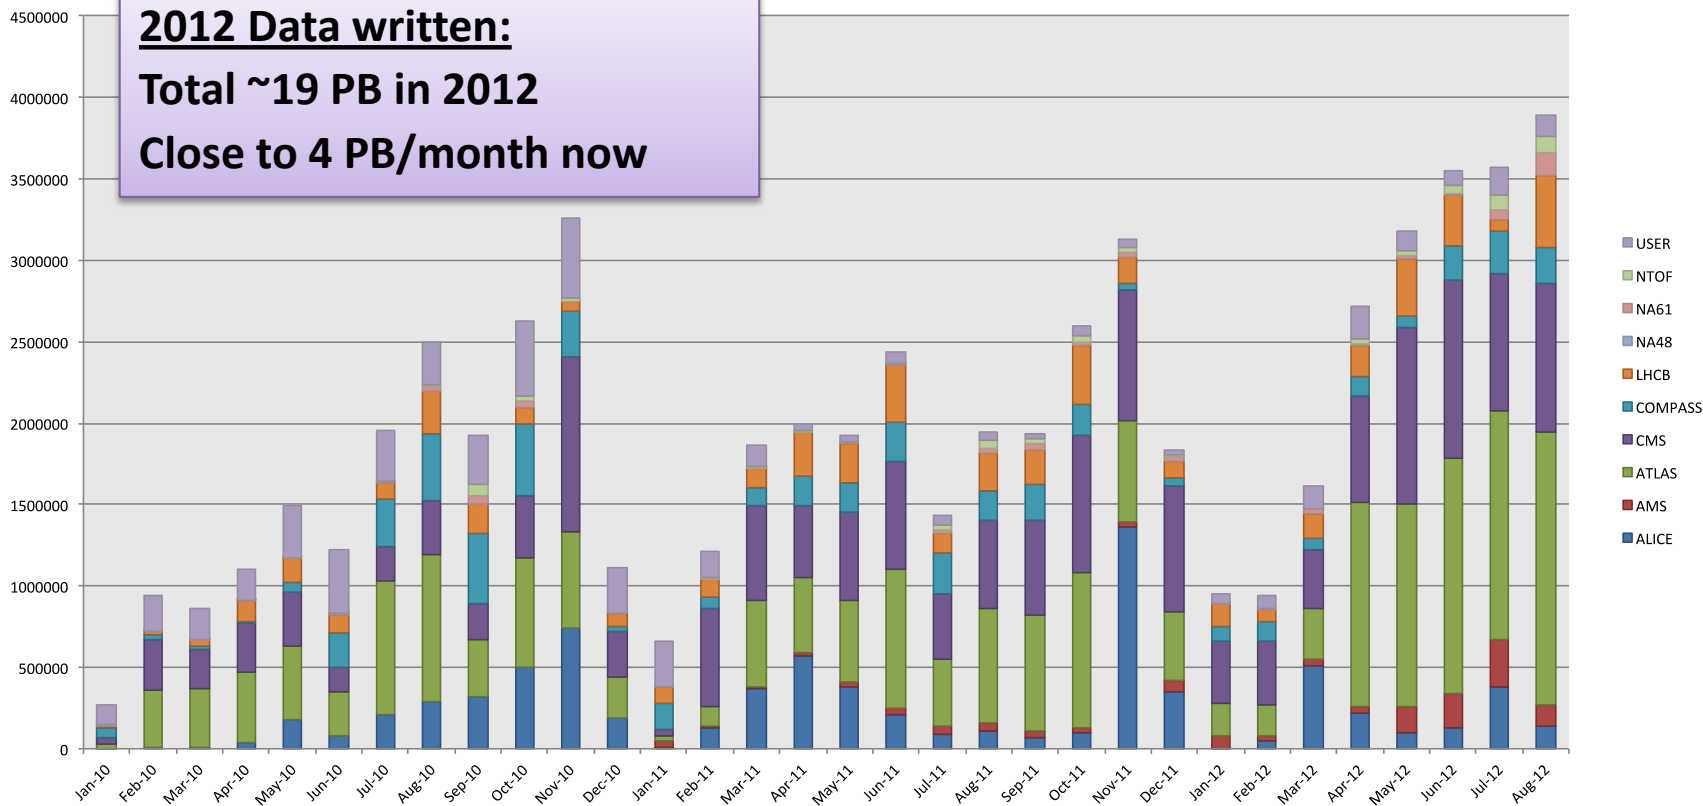
27th September 2012




Castor data written 2010-12

CASTOR data written, 01/01/2010 to 31/8/2012 (in GB)

2012 Data written:
Total ~19 PB in 2012
Close to 4 PB/month now

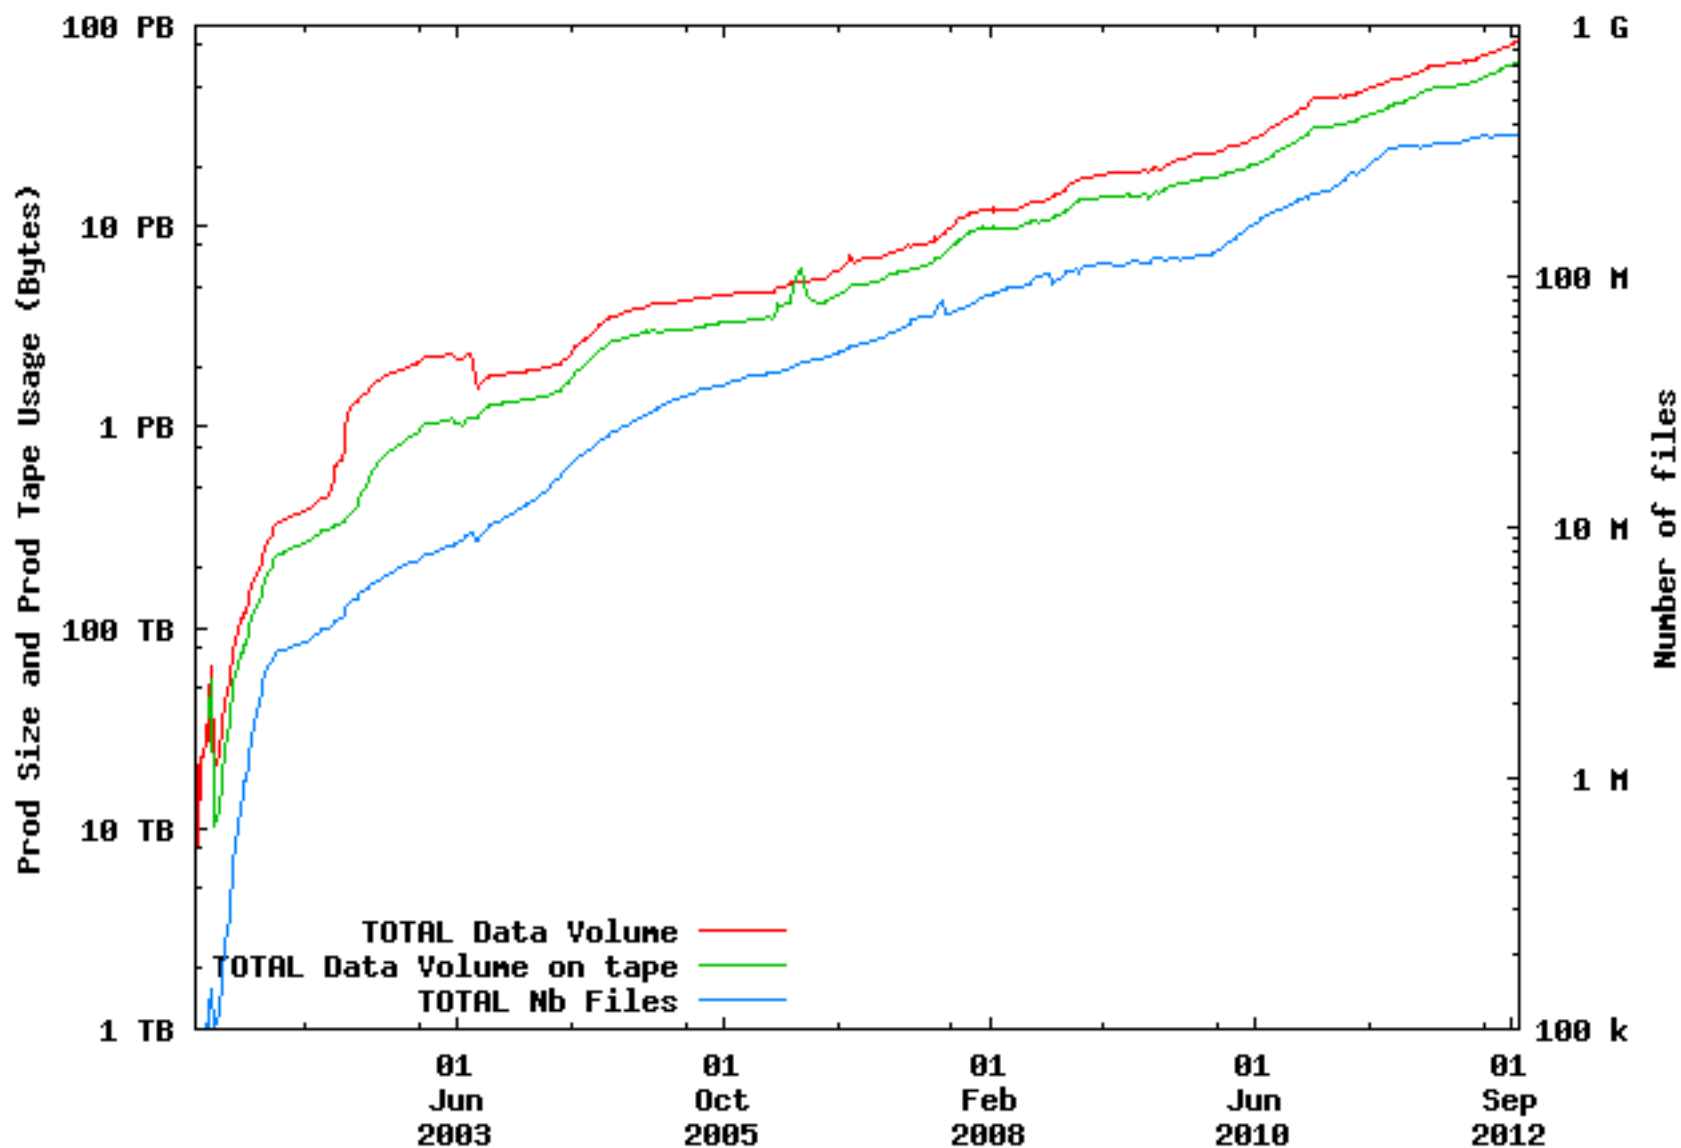


Data rates in Castor increased
 3-4 GB/s input
 ~15 GB/s output





Experiments Production Data in CASTOR




Data in 2012

dashboard

Throughput

2012-06-01 00:00 to 2012-09-14 00:00 UTC




June – Sep 2012

dashboard

Throughput

2012-06-01 00:00 to 2012-09-14 00:00 UTC






A quick glimpse of 2012



- Standard data taking year for ALICE
- p-p – emphasis on rare triggers, high Pt (Calorimeter)
- pilot p-A run (2.5 million events)
- long p-A run in February 2013 (still counts as ‘2012’)
 - Does not require additional computing resources
- Stable performance of all Grid sites, exceptional of T0

Increase of CPU for analysis trains, proportional decrease of chaotic



ATLAS – Tier 1 (and 0)


Current Tier 1 Status




- ATLAS Grid activities at Tier 1s:
 - As stated, we have more CPU as well as (but not that much) disk, we need to be careful!
 - Performing very well!
 - There are two upcoming data reprocessings:
 - 8 TeV pp data,
 - Heavy ion data from 2011.

Tier-1 Disk [PB]	Predicted	Actual	Revised model
RAW data	1.9	1.4	3.5
Real ESD+AOD+DPD	3.8	4.5	4.5
Simulated data	8	7.1	8
Calibration and alignment outputs	0.4	0.4	0.4
Group data	6	3.2	6
User data (scratch)	2	2	2
Cosmic ray data	0.2	0.2	0.2
Processing and I/O buffers	4.3		4.3
Dynamic Data Buffers	Included in WLCG scaling factor	8	Included in WLCG scaling factor
Total	27	26.8	29


Tier-1 CPU [kHS06]	Predicted	Actual	Revised model
Reprocessing	23	1	29
Simulation Production	71	248	71
Simulation Reconstruction	134	109	134
Group+User activities	60	78	60
Total	289	436	295



■ MC Simulation ■ MC Reconstruction ■ User Analysis ■ Group Production ■ Others
■ Data Processing




■ mc11 ■ detail1
■ data10 ■ user



2012 data types breakdown, to CERN tape

Total 7 PB to date
ESD 4.3 PB

RAW 1.7 PB
AOD 0.5 PB
Other 0.5 PB



ATLAS – Tier 2

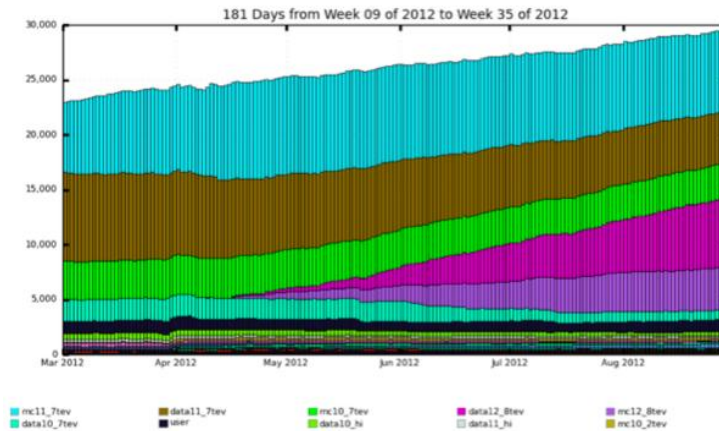

Current Tier-2s Status



- Tier-2s:
 - Successful running of digitization and reconstruction, implemented beginning of the year.
 - Now also taking a big part in group (D3PD/NTUP) production.
 - Repositories for all 'active' data and MC.
 - About 30% of disk space used for 'dynamic data buffers':
 - Popular data (group production, AODs..) replicated dynamically by PD2P.

Tier-2 CPU [kHS06]	Predicted	Actual (average Mar-Aug)	Revised model
Simulation Production	199	395	199
Simulation Reconstruction	29	79	29
Group + User activities	90	138	90
Total	319	612	319

Tier-2 Disk [PB]	Predicted	Actual	Revised model
Real AOD+DPD	11	5.4	13
Simulated data	22	11	22
Calibration and alignment output	0.3	0.3	0.3
Group data	10	4.8	10
User data	2	1.4	2
Processing and I/O buffers	1		1
Dynamic Data Buffers	Included in WLCG scaling factor	7.1	Included in WLCG scaling factor
Total	47	30	48




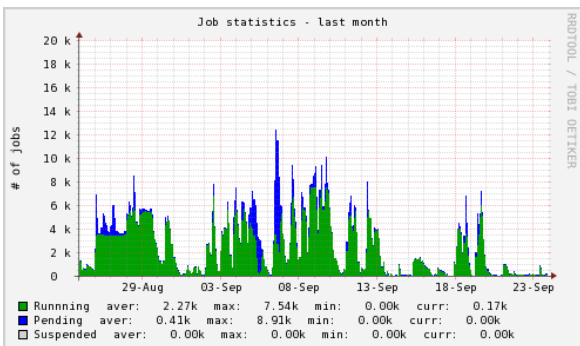
CMS

Tier-0 is working well

- CPU efficiency is between 80-90%
- Tier-0 can spill over to the public queues
 - Used mostly to catch up after long or particularly high luminosity runs

Simulation Production is roughly what we anticipated



- ~400M simulated events
- ~1B reconstructed events



Utilization of Tier-1s is high

- Averaging 102% of pledged capacity used over the last 12 months

Activity remains high for analysis



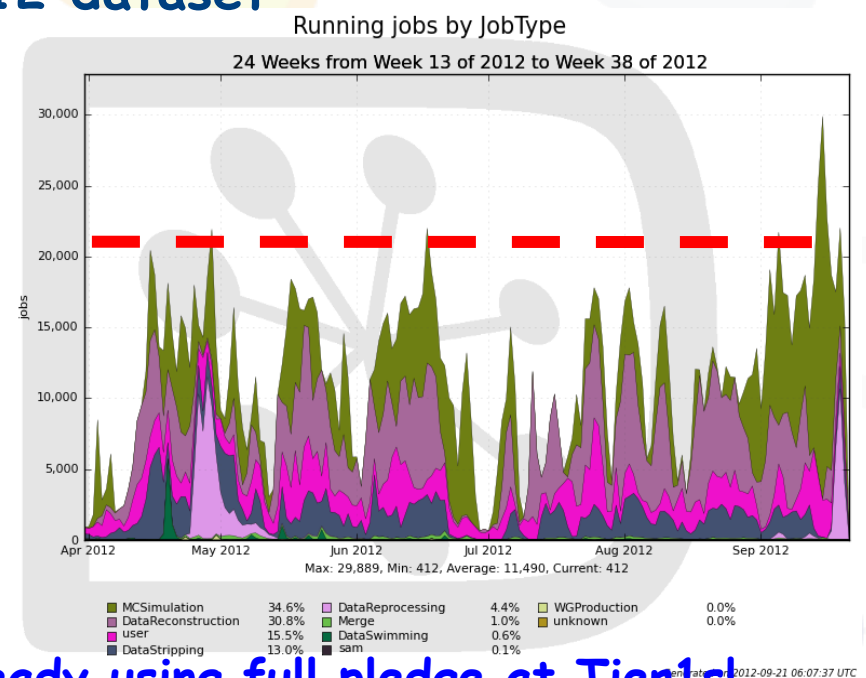
Bird@cern.ch





LHCb computing production activities in 2012

- Prompt processing (reconstruction and stripping) of new data
- "Swimming" of 2011 stripped data
- MC production
 - Large samples for analysis of 2011 data
 - Preliminary samples for 2012 data
- Reprocessing of complete 2012 dataset
 - Started 17th September
- Resources OK until now:
 - CPU usage 50% of pledge
 - ↳ As expected for first half of year
 - ↳ now ramping to >100% with reprocessing
 - 2012 disk pledge ~sufficient also for reprocessed data
 - ↳ Active data management
 - Major shortage of tape, already using full pledge at Tier1s!
 - ↳ Pledge based on 2012 request, assuming shorter run
 - ↳ Major problem for the future (rest of 2012, but also 2013)





Storage: forecast for March 2013

- Disk OK but big imbalance between sites (backup slide)

Disk	Pledge	2012	
	PB	PB	%
Tier0	3.5	3.5	36
Tier1	7.3	6.3	64

- Serious shortfall in tape, no solution yet
 - Already cleaned up all 2012 prompt SDST to fit in existing tape
 - No room at Tier1s for second copy of new RAW, and for ONLY copy of FULL.DST

Current Tape (PB)	Tier0	Tier1
RAW	2296	2020
SDST	152	1025
Archive	1544	2251
Total Current	3992	5296
Tape March 2013 (PB)		
RAW	3096	2820
FULL.DST	1350	4449
Archive	1675	2775
Total March 2013	6121	10044
Pledges	6400	5324

WLCG Operations

- Operations over the summer quite smooth
- No particular problems
- Resource usage consistently high

Tier 0 upgrades

- CERN CC – extension
 - Scheduled for completion Nov 2012 – still on track
 - Required for 2013 equipment installation
- Wigner centre
 - Site visit last week – progress on schedule
 - Expect to be able to test first installations in 2013
 - Networking CERN-Wigner (2x100 Gb): procurement ongoing
 - Latency testing for 2 months
 - Fraction of Ixbatch with 35 ms delay – no observed effects

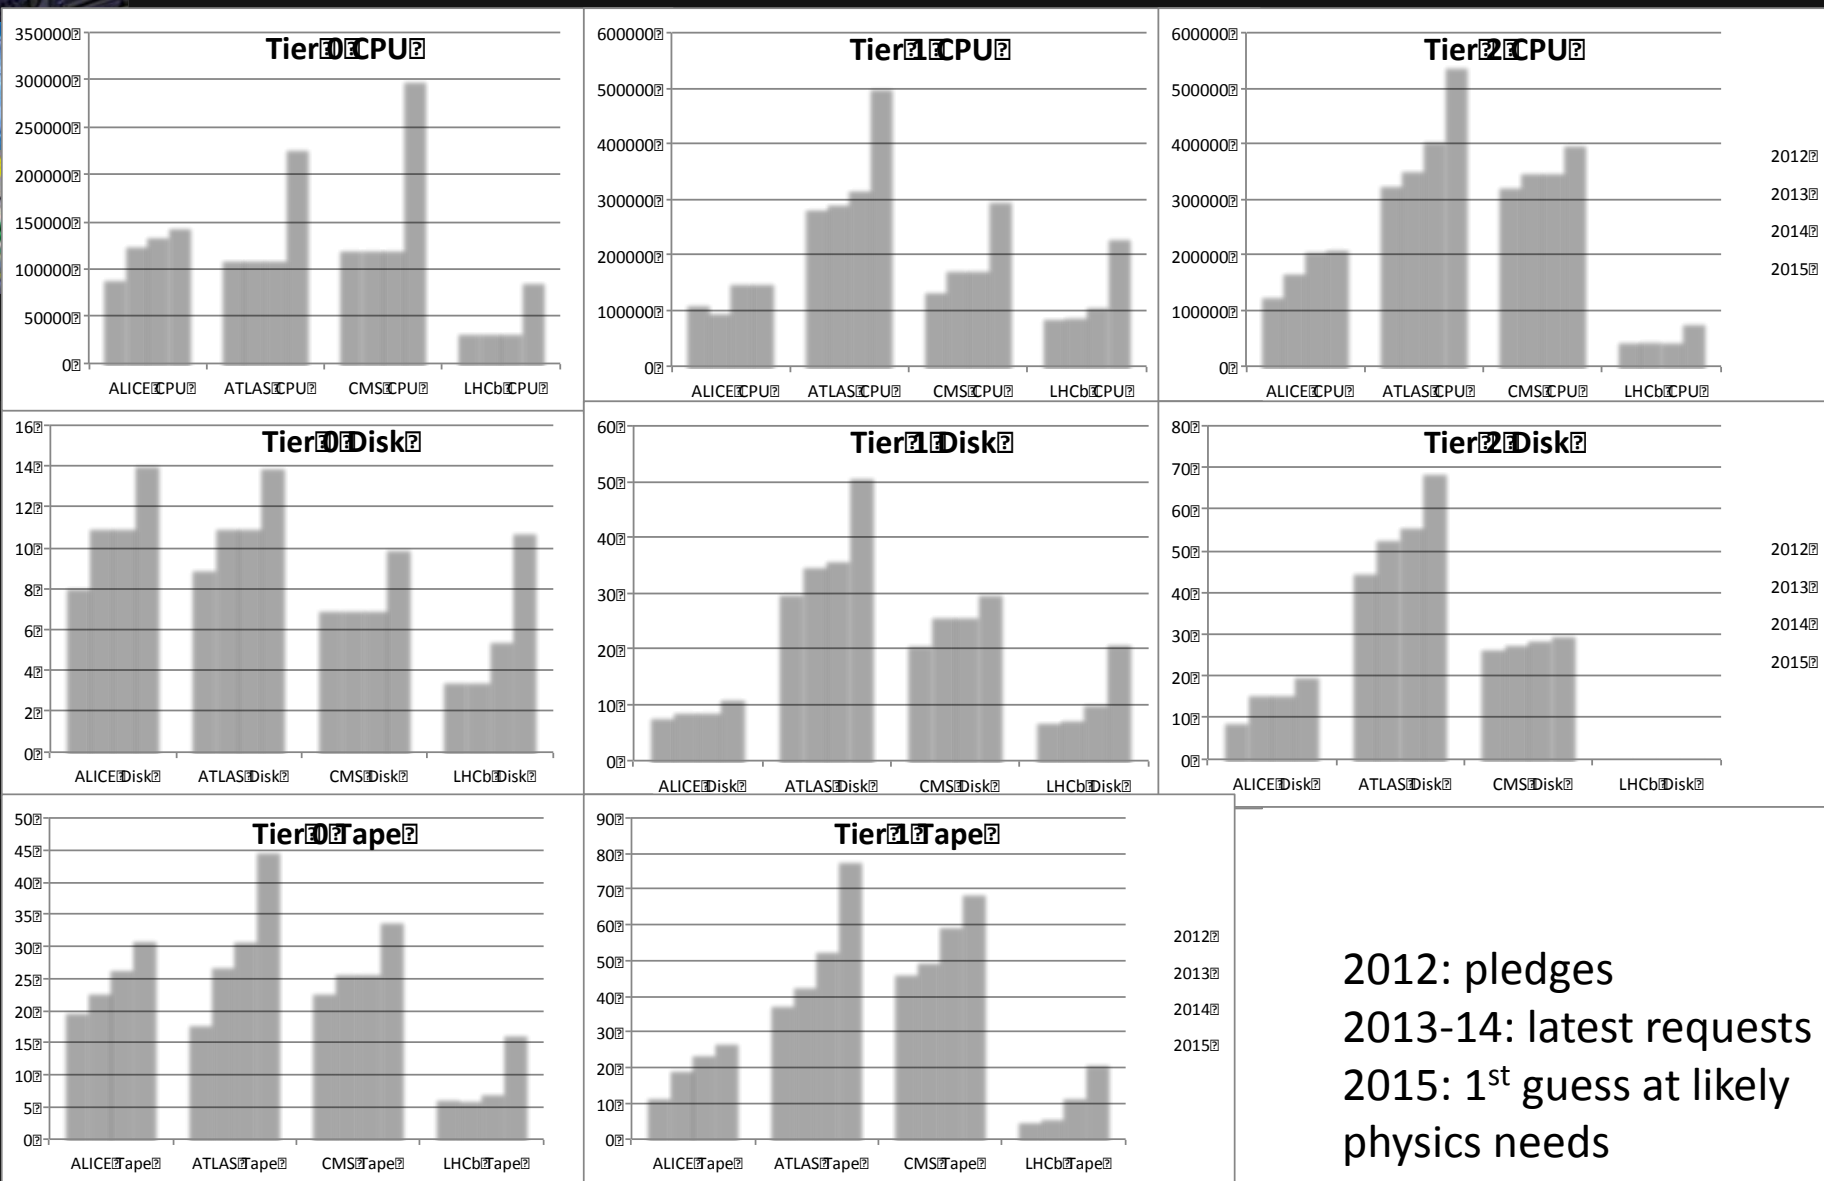
Resource evolution

- 2012
 - Extended run 2012-2013 has implications for 2012 resources: ~20% additional resources
 - In most cases the hope for additional 2012 resource at the end of 2012 cannot be satisfied
- 2013-14
 - Requests for 2013 revised to take into account the effects of the extended 2012 run
 - 2014 requests close to the 2013 revised requests – some slight increases needed for analysis work and simulation
 - Computing activities in LS1:
 - Full re-processings of complete 2010-12 data
 - Simulations needed for 2015 at higher energy
 - Analysis...

First look at 2015

- Statement of 2015 needs is not yet formally required, but
- Important to send message to RRB that continued investment in computing essential to fully exploit the performance of the LHC
- Experiments have made first rough estimates of potential requirements for 2015 based on anticipated physics needs, and likely LHC running conditions
 - ALICE: standard data taking year – increases driven only by data volumes ($\sim 20\%/yr$)
 - ATLAS: possible goal is 1 kHz trigger rate; realise would have to work on software and computing models to achieve within reasonable resources
 - CMS: Could also sustain 1 kHz trigger rate; would give x3 events/time
 - LHCb: max possible is 700 MB/s from DAQ (double 2012) – still many unknowns in desired trigger rates; this would imply x2 resources cf 2014

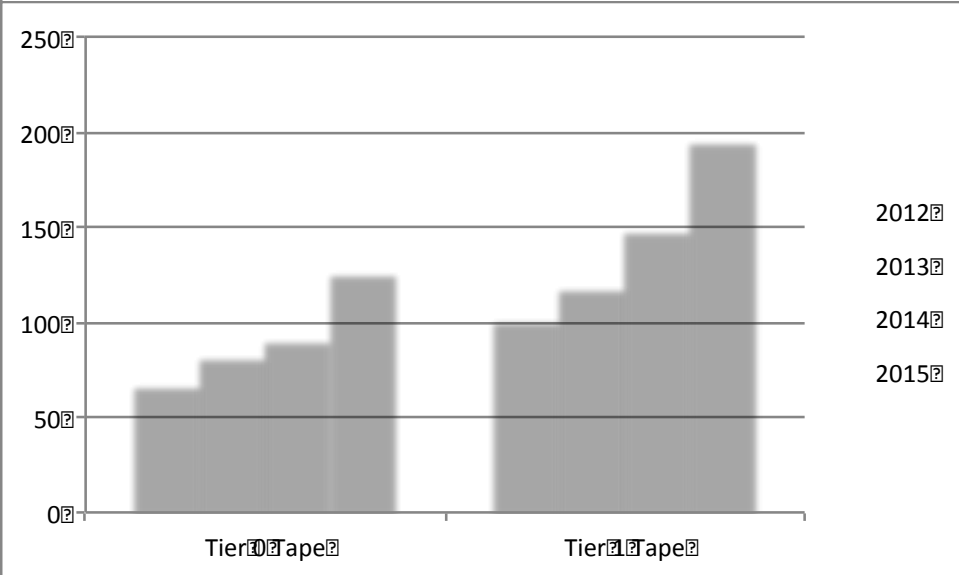
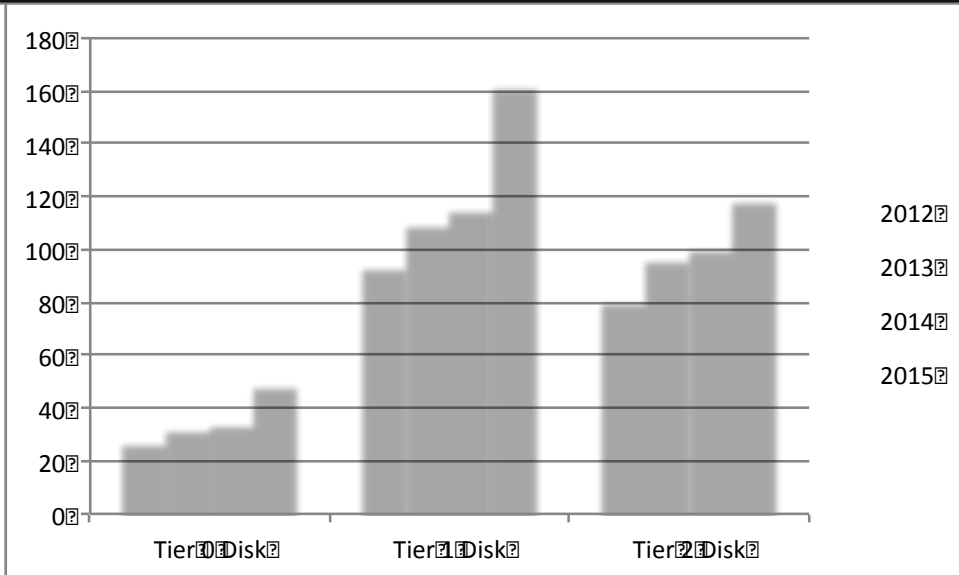
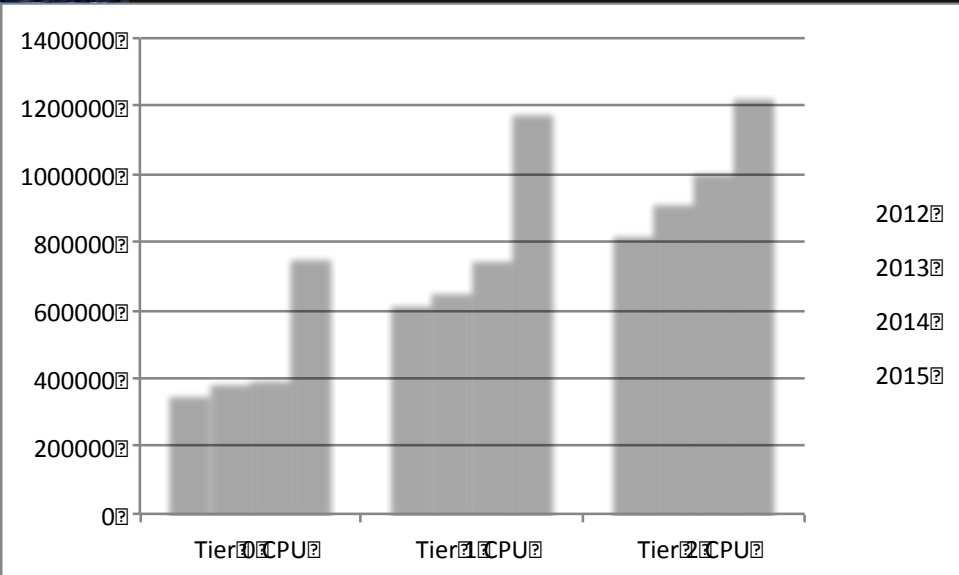
Evolution of resources



2012: pledges
 2013-14: latest requests
 2015: 1st guess at likely physics needs



Overall evolution



NB: while 2015 looks high, not so different from 25% yearly increase from 2012 (which is ~flat budget)



Plans for LS1



- A selection of topics:
- Improvements/simplification of the Event Data Model (EDM)
 - Some of the C++ classes designed at the time are not well suited for modern CPU architectures
 - Data components scattered over memory: bad for memory cache usage
 - And bad for vector function usage: compilers can't auto-vectorize
 - ⇒ Design simplified EDM, for usage within algorithms; start with tracking algorithms
- Concurrency: usage of many-core designs with small memory footprint
 - Multi-core used since a while with multi-process software (not multi-threaded); memory footprint reduced 10-1/3, still too big for many-core
 - Multi-threaded designs being tackled algorithm by algorithm
 - Need some support in the software framework: underway
- benefitting from collaboration with IT/OpenLab, PH-SFT and CMS in these and related fields
- New Distributed Data Management system (DDM) being implemented
- WAN data access and data caching (file level, event level)
- Applications-driven usage of networks
- New MC production system
- Simulation: Integrated Simulation Framework (ISF)
 - Selection of full/fast/parametrized mode per subevent
 - Concurrency within event
 - Plus benefit from reconstruction speedup

9

Investment of effort on significant software improvements, adaptation of computing models, etc. to maximise efficiencies of use of available resources is also essential (and some efforts are underway)