



CMS

Computing Report to LHCC referees

- **Production summary, CSA07++**
- **PADA Taskforce**
- **CCRC08-February tests**



Compact Muon Solenoid

Matthias Kasemann
LHCC February 2008



Production Summary 10'07-02'08

- 160M Monte Carlo events produced since October 07
 - On request of Physics, DPG and HLT groups
- Total CSA07 event counts:
 - 80M GEN-SIM
 - 80M DIGI-RAW
 - 80M HLT
 - 330M RECO
 - 250M AOD
 - 100M skims (mixed RECO/AOD)
 -
 - 920M events
- Events were processed + reconstructed in several steps, several times

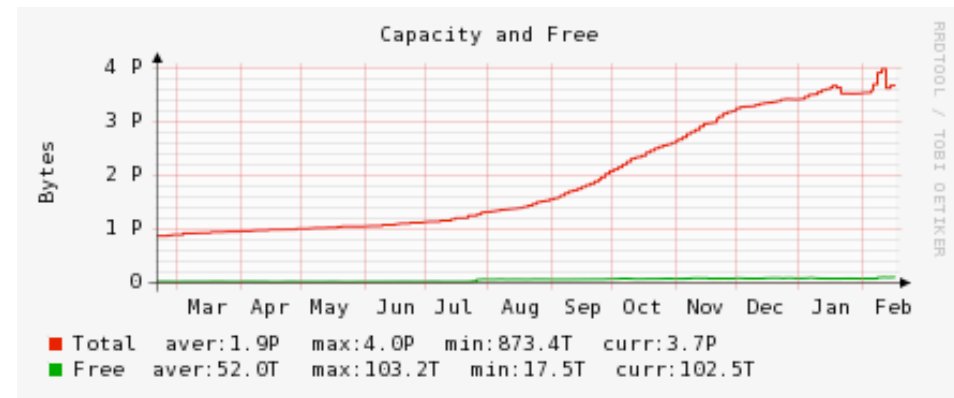


Total: 477797046.00 , Average Rate: 39.50 /s



CSA07 event samples

- **CSA07 signals:**
80M GEN-SIM, 80M DIGI-RAW, 80M HLT, 80M RECO
- **CSA07 soups:** 250M RECO, 250M AOD, 100M skims
 - Three calibrations applied: 10/pb⁻¹, 100/pb⁻¹ 0/pb⁻¹
 - Events produced: RECO, EXPRESS, AOD, skim, ALCARECO
- The CSA07 signal samples really evolved over time. We started from 50M and went up to 85M by now (not a real problem)
- Data volume of CSA07 samples right now (without counting repetitions): 1.9 PB
- Delivery of the samples is mostly done with small remainders pending.

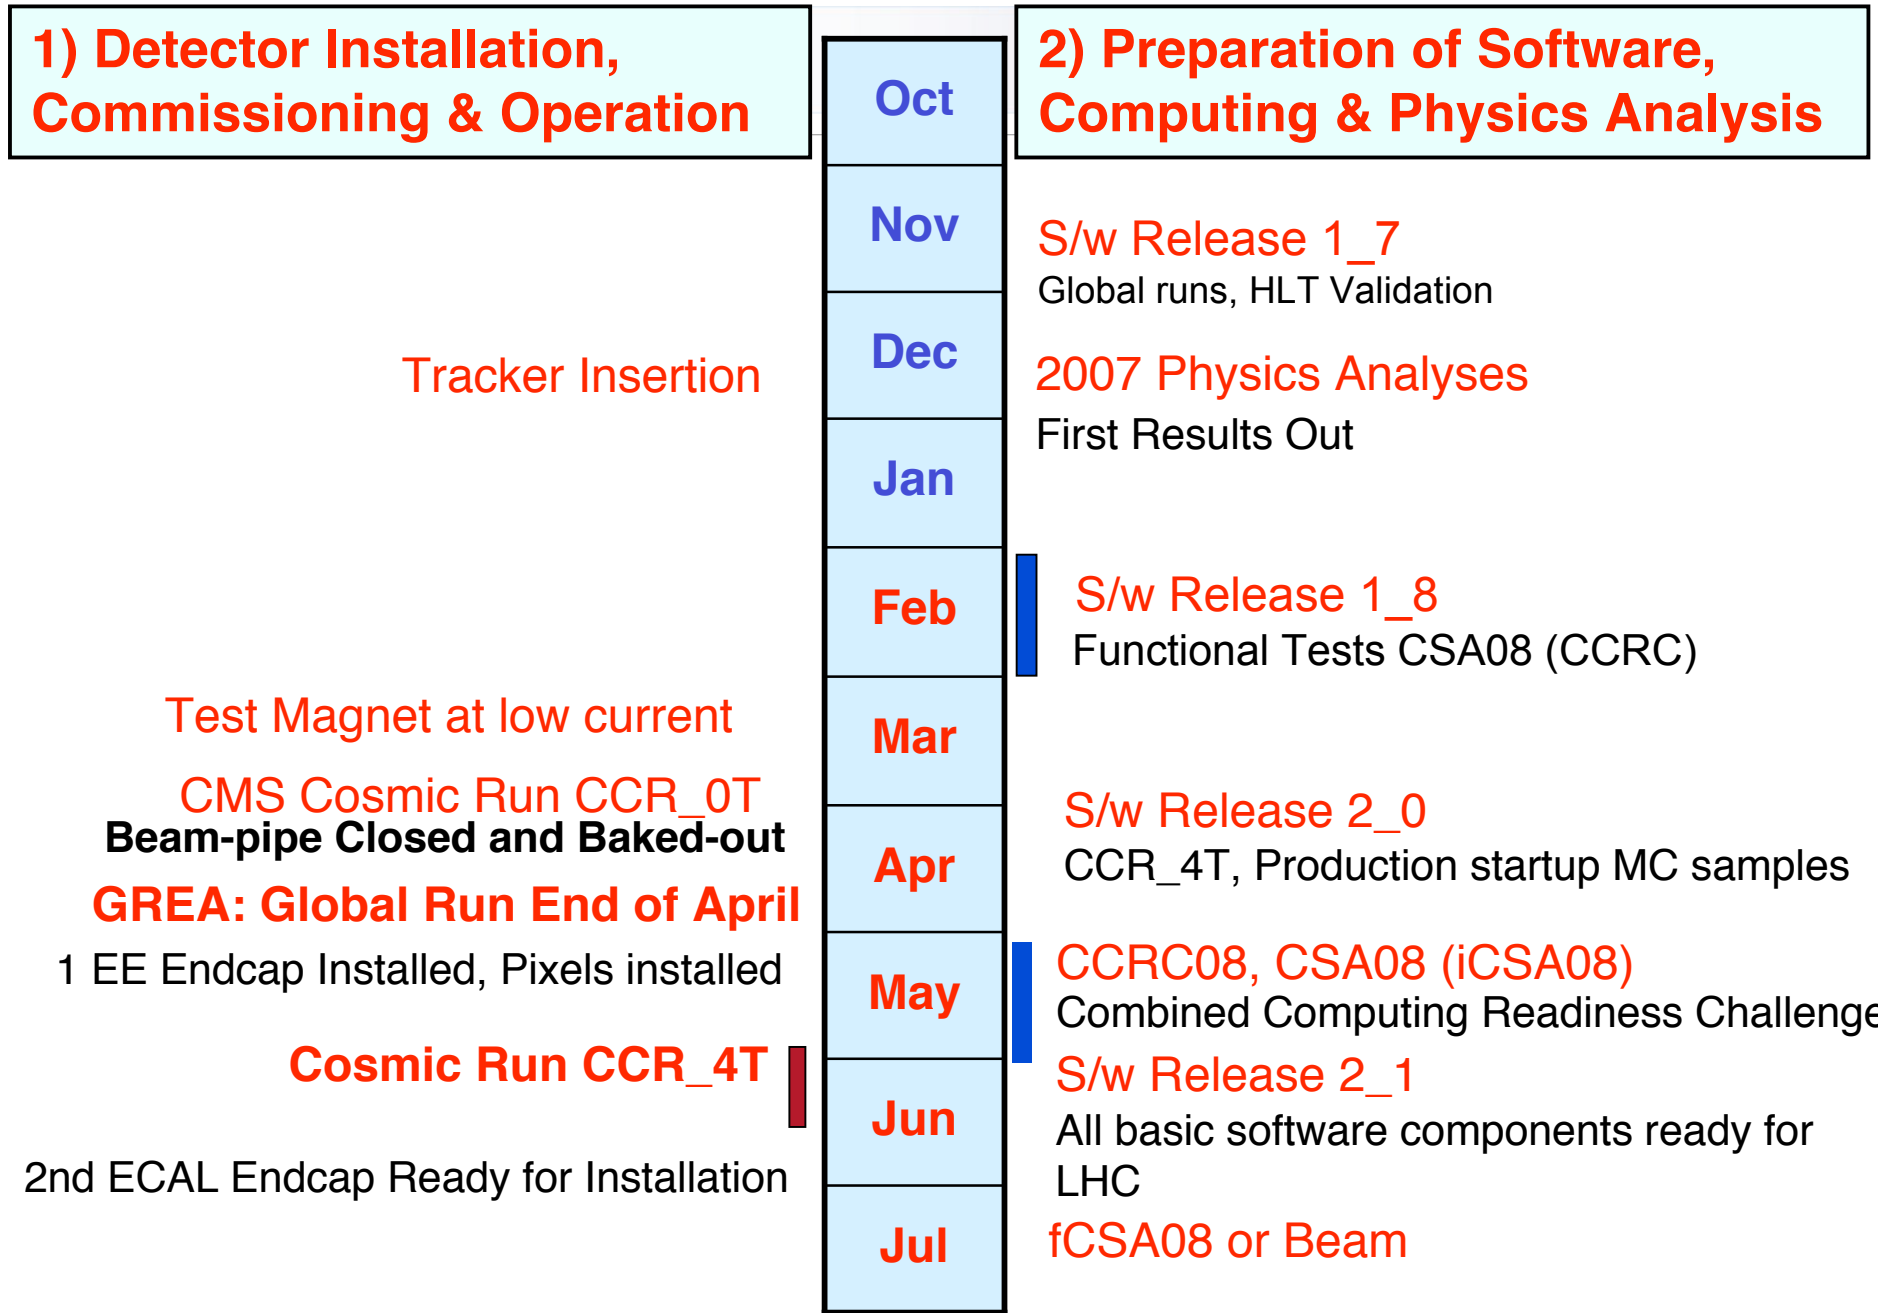


CMSA data in CASTOR@CERN: 3.7PB



CSA07 Analysis Summary

- In CSA07 a lot was learned and a lot was achieved...
- The production infrastructure is in full operations
- CSA07 analysis identified tasks to be addressed
 - For Offline: detailed list of development work
 - Two strategies for Computing:
 - Integrating development, deployment and commissioning in a new task force:
Processing And Data Access (PADA)
- coordinated by I.Fisk and J.Hernandez
 - Testing the computing infrastructure in **CCRC08/CSA08** in February and May '08





Processing and Data Access: PADA

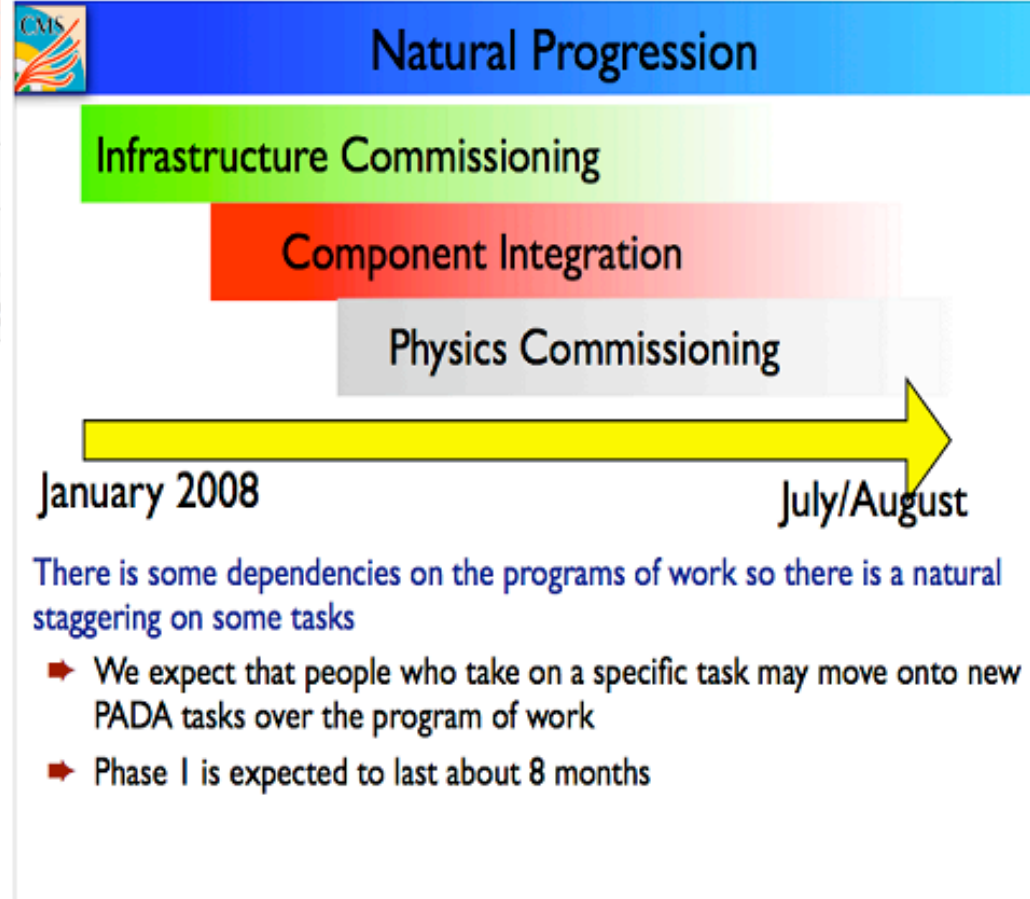
- **The Processing and Data Access Task Force is a series of tasks and programs of work**
 - designed to bring the elements of the Computing Program into stable and scalable operations.
- **It covers the integration, commissioning and scale testing of data processing workflows at the Tier-1 and Tier-2 sites and includes:**
 - Validation of the infrastructure for organized processing and user analysis
 - including the sites and the CMS workload and data management tools.
 - Validation of the distributed production system
 - performing functionality, reliability and scale tests.
 - Help sites to commission, configure and optimize the networking and storage through scale testing data transfers and data processing.
 - Improvement in the efficiency of accessing data across the CMS computing system from global transfers to local access.



PADA tasks + schedule

Transfer Commissioning (DDT Phase 2)
Site Commissioning
CMS Service Testbed
Production Component Validation
Analysis Server Validation
Monitoring and Information Integration
Dynamic Tier-2 Data Management
User Driven Organized Processing
Distributed Analysis Functionality and Scale Testing

- Succeeded to find names for 5 tasks, more to go





CCRC08 February tests

Data recording at CERN

1a) readout from P5, use HLT, w. stream definition, use Storage Manager, transfer to T0, perform repacking, write to CASTOR (D.Hufnagel)

- **Goal: verify dataflow for CMS, commission the new 10GB fiber**
- **Status:**
 - 13.2.08: First successful transfer on new 10 GB fibre at 100MB/s (limited by transfer node)
 - Next step: integrate into transfer system run in parallel to normal data transfers




CCRC08 test: Data recording at CERN

1b) CASTOR data archiving test (M.Miller / DataOps team)

- Goal: verify CASTOR performance at full CMS and ATLAS rate
- Status: very successfully completed, reached rate of 1.5 GB/s
 - Good coordination with CERN-IT, quick response
 - Test at all-VO rate, other VO's didn't stress the system

Migration
policy mod.



Recovered

Ran out of
tapes

t0export: in from WN, out to tape



CCRC08 test: High Rate Processing at T0

Coordinated by M.Miller / DataOps

Goal:

- “high-rate” processing of cpu/RAM limited jobs
- Originally: measure interaction with other VO’s on same WN
BUT: CMS does not share WN with other VO’s @ CERN (for now)

Setup:

- regular operations (physics requests)
- ReReco with 0pb^{-1} conditions of Stew and Gumbo

Status:

- started with 41k jobs of the 80 TB Stew AllEvents
- Finished in expected time
- Not much action from other VO’s, no sign of WN problems
- Again turning into a CASTOR I/O test


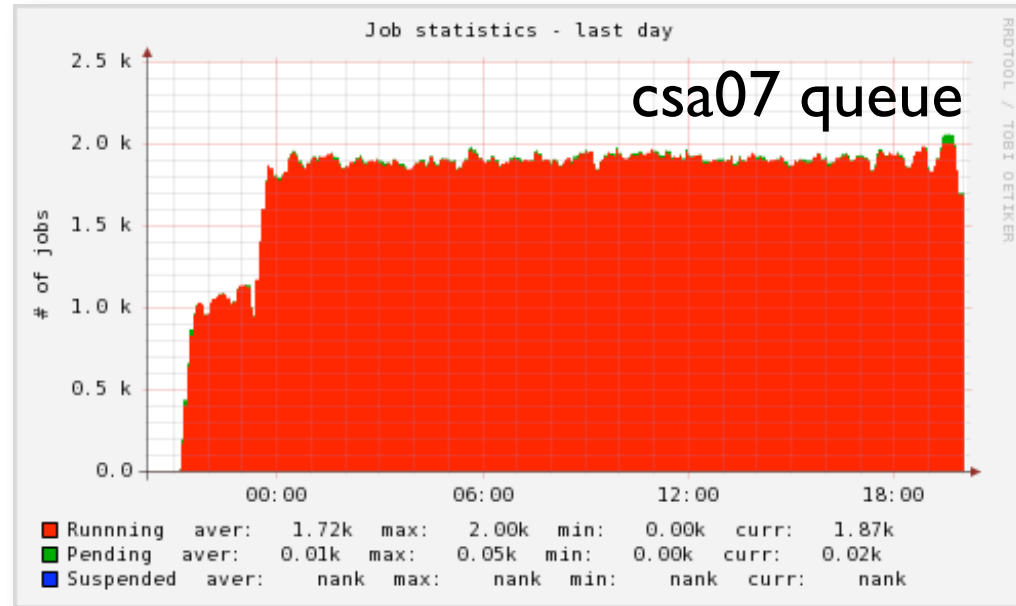


CCRC08 test: High Rate Processing at T0

Wednesday snapshot

Summary:

- processing runs routinely
- Small level of IO errors (2%), cured by retries
- Will test: copy files to local disk





CCRC08 Transfer tests

Goal: use SRMv2 data transfers where possible

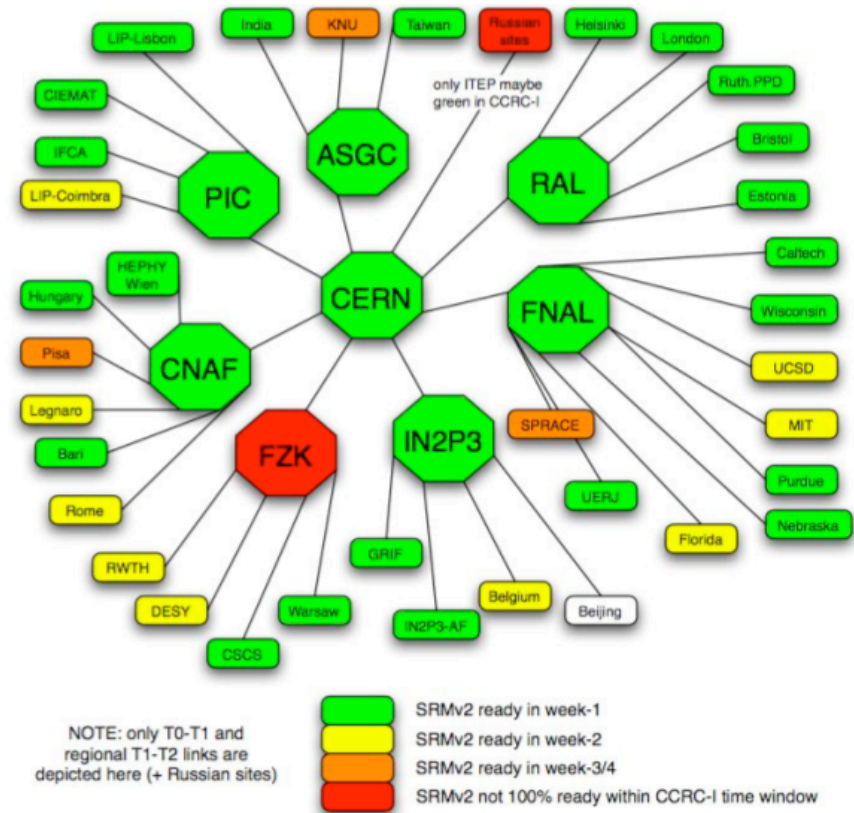
Target rates:

- T0-T1: 25/40/50% of full 2008
- T1-T1: 50% in+outbound
- T1-regional-T2: full/high rate
- T2-regional-T1: full/high rate

A detailed Plan worked out:

- cycle through different parts of all link combinations per week

Status of SRM-v2 deployment (~5.2.08)



Tests are progressing well

- individual problems are being addressed and result in delayed testing
- More detailed analysis available at the end of February



CCRC08 Re-Reconstruction tests

Coordinated by G.Gomez-Ceballos, Josep Flix

Goal: measure performance of:

- **Migration from Tape to Buffer: pre-stage test.**
- **Reprocessing exercise: use all available CMS CPU-slots at T1s**

Plan:

- **Select one (or more) dataset(s) of ~10TB size existing at T1.**
- **Remove all the files from disk (aka, T1_Buffer).**
- **Fire the staging from Tape to Buffer of all files.**
- **Monitor the process and provide some measurements/plots**
- **Run Re-reconstruction over CSA07 data present at all T1s**
 - **Measure performance**




CCRC08 Re-Reconstruction tests

Status:


- buffer to tape migration successfully finished at all sites
 - Results: total staging time 8-44h, rate: 67-250MB/s observed
 - Except IN2P3, performance was poor, reconfigure and redo

ASGC Recall Test: comments and plots

- Recall test data tape throughput:



- Tape->Buffer throughput: **RAL**



- high performance processing without overlap with ATLAS
 - Finished at FNAL(1200 slots), CNAF(1000-1300 slots), FZK(600 slots), ASGC(300 slots)
 - Next to go: IN2P3 and PIC, RAL, but no problem foreseen
- Processing test together with ATLAS planned at two Tier-1's:
 - special queue for Atlas and CMS is setup at IN2P3 and PIC



CCRC08 Monte Carlo tests

Coordinated by DataOps

Goal: Production tests of FastSim Monte Carlo

Status:

- **Physics groups want to use 50M of the CSA07 samples (100pb^{-1} calibration), reading AOD's.**
- **Fast Simulation production based on CMSSW_1_6_9 ready to start**
 - cfg file which can run over RECO and AOD files is tested
 - data ops ready to start the production
(data handling: use resources at T1 or ship AOD data to T2)



CCRC08 CAF tests

Coordinated by P.Kreuzer

Goal:

- ramp-up CAF resources
- verify basics CMS use cases at scale

	CPU	Disk	Tape
T0 2008	70% Dual quad-core (16GB RAM) 30% Dual Dual-core (8GB RAM)		
CAF 2008	1200 slots (128 slots in '07)	1600 TB (35 TB in '07)	3 PB
CAF CCRC'08	250 slots	150-200TB	

note: 3000 slots \approx 5.3MSI2K , 1200 slots \approx 2.1MSI2K

Status: good progress made

- resources configured according to plan
- Regular CAF meetings with user representatives (Global Run, ALCA and Physics)
- Plan for CCRC08 (week 3 and 4):
 - Transfer GR data from T0 to CAF and populate local DBS
 - Finalize RPC workflow
 - Test/Run HcalCallsoTracks workflow
 - Test/Run Muon Alignment workflow
 - Setup and test CRAB, local submission
 - Collect list of CAF groups and users per group. Provide to IT, both for batch/interactive CAF



Computing Summary

- The Computing infrastructure is fully utilized for ongoing production
 - Finished original CSA07 production (and much more)
- Detailed analysis of CSA07 performance was performed.
 - Direct result for Computing:
defined PADA tasks and CCRC08 functional tests
- The PADA taskforce addresses deployment, integration, commissioning and scale testing. It will bring the elements of the Computing Program into stable and scalable operations.
- The CCRC08 functional tests in February complements CSA07 and test additional functionality:
 - Important scale testing at T0
 - Data-handling and full-scale reprocessing at T1
 - Data transfer tests using the new SRMv2 protocol
 - Role-out of Fast Simulation for Monte Carlo Production
 - Ramp-up of CAF resources and scale testing of CAF use-cases
- Detailed planning of CCRC08(May), iCSA08 and f-CSA08 is going on, expect to agree on initial scope and goal during CMS week.


The image features the CMS logo in a red serif font at the top. Below it, the text 'Backup slides' is written in a bold blue sans-serif font. The background consists of several light blue, curved lines that sweep across the frame, creating a sense of motion or a stylized globe.

CMS

Backup slides



Last 2 weeks: integrated tape system usage



Aggregate [MB/sec]	rfcp mean [MB/sec]	rfcp sigma [MB/sec]	#jobs	t _{lumi} [sec]
800	52.5	17.4	~30	37
1100	40.5	15.4	~50	27
1300	31.7	13.0	~90	22
1500	18.1	11.5	100-infinity	18

Rates ultimately limited by 1 Gbs on 13 t0input servers

CMS CASTOR TEST - Performance observed:

Averaged 633 MB/sec write (1.1 GB/sec during test)

Averaged 548 MB/sec read (~400 MB/sec during test)

Read Spike: regular stagein, 101 drives => 5 GB/sec