

# Expt SC3 Status

Nick Brook

In chronological order:

ALICE


CMS

LHCb

ATLAS




# Alice Physics Data Challenge '05 - goals




- PDC'05 : Test and validation of the remaining parts of the ALICE Offline computing model:
  - Quasi-online reconstruction of RAW data at CERN (T0), without calibration
  - Synchronised data replication from CERN to T1's
  - Synchronised data replication from T2's to their 'host' T1
  - Second phase (delayed) reconstruction at T1's *with calibration* and remote storage
  - Data analysis
- Data production:
  - List of physics signals defined by the ALICE Physics Working Groups
  - Data used for detector and physics studies
  - Approximately 500K Pb+Pb events with different physics content, 1M p+p events, 80TB production data and few TB user generated data
  - Structure – divided in three phases:
    - Phase 1 – Production of events on the GRID, storage at CERN and at T2s.
    - Phase 2 ( synchronized with SC3) – Pass 1 reconstruction at CERN, push data from CERN to T1's, Pass 2 reconstruction at T1s with calibration and storage:
      - Phase 2 (throughput phase of SC3) – how fast we can push data out
    - Phase 3 – Analysis of data (batch) and *interactive analysis with PROOF*

# Methods of operation




- Use LCG/EGEE SC3 baseline services:
  - Workload management
  - Reliable file transfer (FTS)
  - Local File Catalogue (LFC)
  - Storage (SRM), CASTOR2
- Production and data replication phases are synchronised with LCG/EGEE SC3
- Operation of PDC'05/SC3 coordinated through ALICE-LCG Task Force
- Run entirely on LCG resources:
  - Use the framework of VO-boxes provided at the sites
- Require approximately 1400 CPUs (but would like to have as much as possible) and 80 TB of storage capacity
- List of active SC3 sites for ALICE:
  - T1's: CCIN2P3, CERN, CNAF, GridKa (up to few hundred CPUs)
  - T2's: Bari, Catania, GSI, JINR, ITEP, Torino (up to hundred CPUs)
  - US (OSG), Nordic (NDGF) and a number of other sites joining the exercise presently
  - SC3 + others – approximately 25 centres




# Job submission

VO-Box  
LCG  
User Job  
ALICE catalogues




# Status of production



- Setup and operational status of VO-boxes framework:
  - Gained very good experience during the installation and operation
  - Interaction between the ALICE-specific agents and LCG services is robust
  - The VO-box model is scaling with the increasing load
  - In production since almost 1 ½ months
- Many thanks to the IT/GD group for the help with the installation and operation
- And to the site administrators for making the VO-boxes available
- Setup and status of storage:
  - ALICE is now completely migrated to CASTOR2@CERN
  - Currently stored 200K files (Root ZIP archives), 20TB, adding ~4K files/day
- Operational issues discussed regularly with IT/FIO group, ALICE is providing feedback


# Status of production



- Current Job status:
  - Production job duration: 8 ½ hours on 1KSi2K CPU, output archive size: 1 GB (consists of 20 files); Total CPU work: 80 MSi2K hours; Total storage: 20 TB



# ALICE plans:



- File replication with FTS:
  - FTS endpoints tested at all ALICE SC3 sites
  - Start data migration in about 10 days, initially T0->T1
  - Test, if possible, migration Tx->Ty
- Re-processing of data with calibration at T0/T1:
  - AliRoot framework ready, currently calibration and alignment algorithms implemented by the ALICE detector experts
  - Aiming for GRID tests at the end of 2005
- Analysis of produced data:
  - Analysis framework developed by ARDA
  - Aiming at first controlled tests beginning of 2006





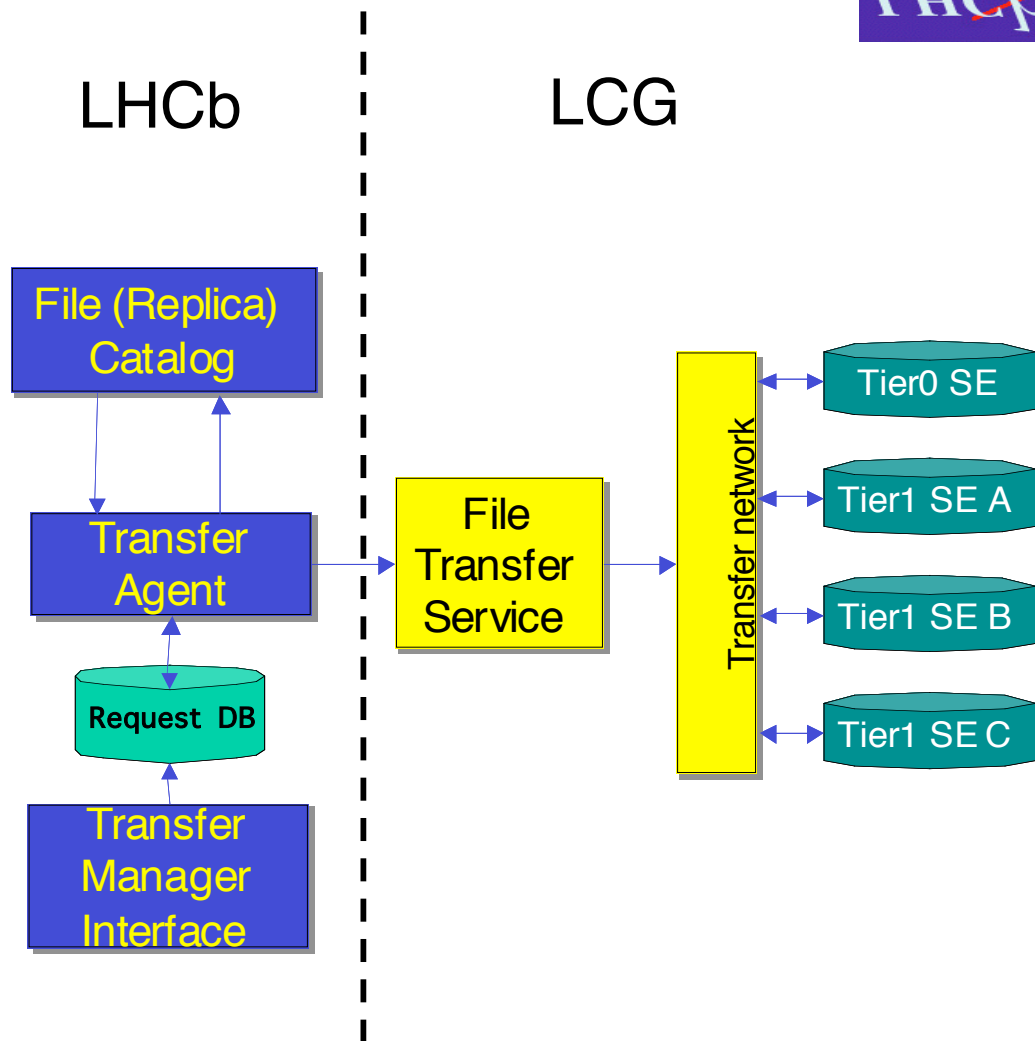
## SC3 Aims

- **Phase 1: (Data Moving)**
  - Demonstrate Data Management to meet the requirements of the Computing Model
  - Planned: October-November
- **Phase 2: (Data Processing)**
  - Demonstrate the full data processing sequence in real time
  - Demonstrate full integration of the Data and Workload Management subsystems
  - Planned: mid-November + December

Currently still in Phase 1 - Phase 2 to start soon

# LHCb Architecture for using FTS

- Central Data Movement model based at CERN.
  - FTS+TransferAgent+RequestDB
- TransferAgent+ReqDB developed for this purpose.
- Transfer Agent run on LHCb managed lxgate class machine



- Gets transfer requests from Transfer Manager
- Maintains the pending transfer queue
- Validates transfer requests
- Submits transfers to the FTS
- Follows the transfers execution, resubmits if necessary
- Sends progress reports to the monitoring system
- Updates the replica information in the File Catalog;
- Accounting for the transfers
  - <http://fpegaes1.usc.es/dmon/DIRAC/joblist.html>

## Phase 1

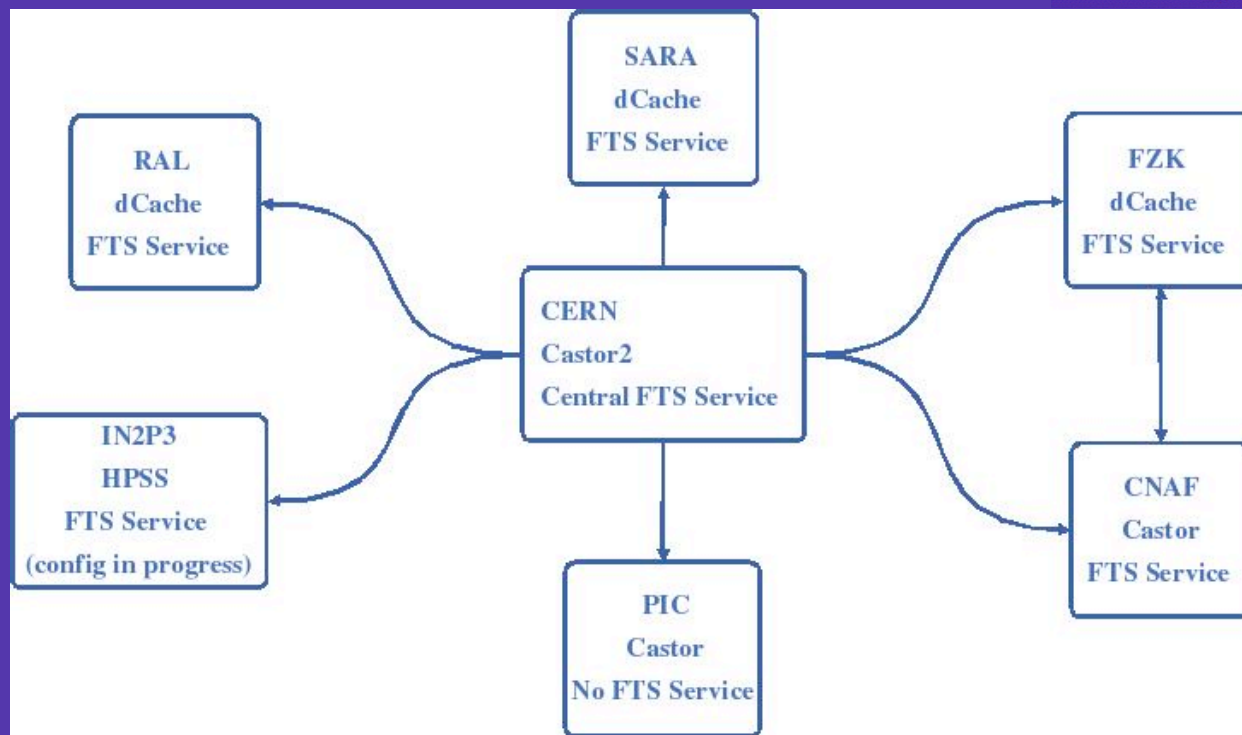
- Distribute stripped data Tier0 → Tier1's (1-week). 1TB
    - The goal is to demonstrate the basic tools
      - Precursor activity to eventual distributed analysis
  - Distribute data Tier0 → Tier1's (2-week). 8TB
    - The data are already accumulated at CERN
    - The data are moved to Tier1 centres in parallel.
    - The goal is to demonstrate automatic tools for data moving and bookkeeping and to achieve a reasonable performance of the transfer operations
- Removal of replicas (via LFN) from all Tier-1's
  - Tier1 centre(s) to Tier0 and to other participating Tier1 centers
    - data are already accumulated
    - data are moved to Tier1 centres in parallel
    - Goal to meet transfer need during stripping process

## Participating Sites

Tier0-Tier1 channels  
over dedicated  
network links

Bi-directional FZK-  
CNAF channel on  
open network

Tier1-Tier1 channel  
matrix requested  
from all sites - still  
in the process of  
configuration -  
central or expt  
coordination ?



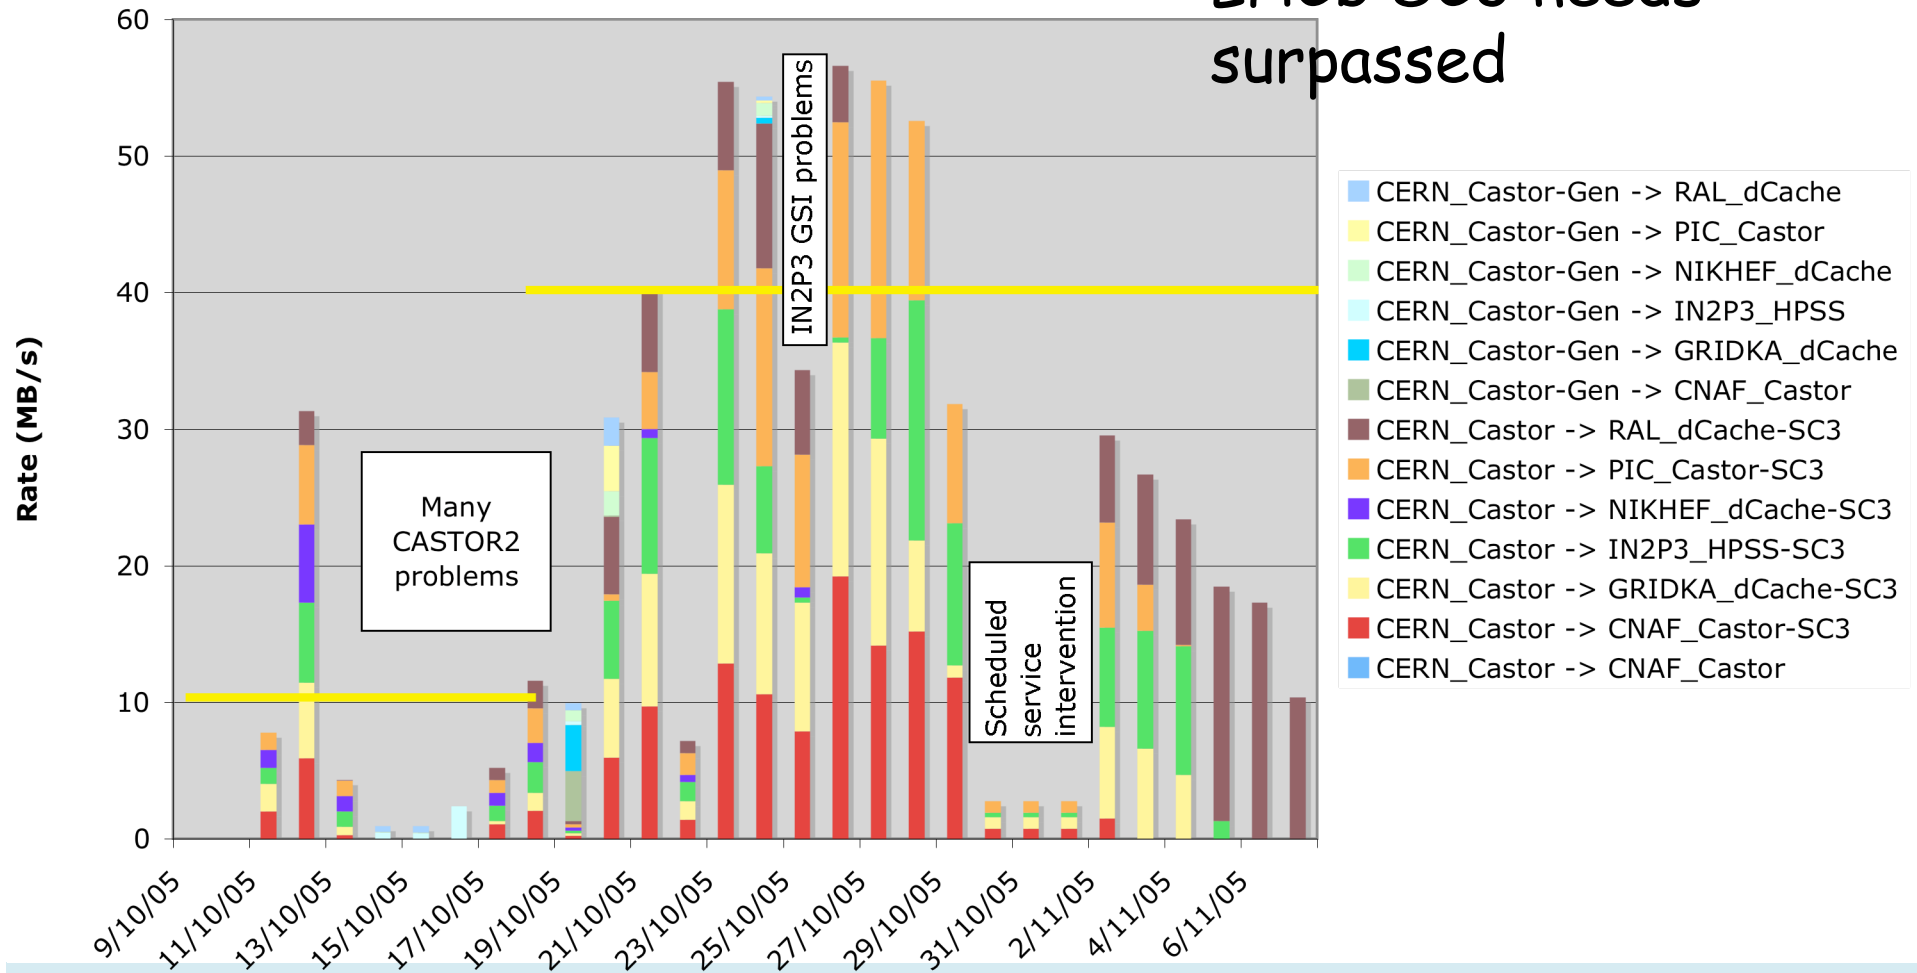
Need for central service for managing  
T1-T1 matrix ??

# Overview of SC3 activity

SARA show almost no effective bandwidth from 25/10

LHCb SC3 Activity

When service stable - LHCb SC3 needs surpassed



## Problems...

FTS files per channel dramatically effects performance

- By default set to 30 concurrent files per channel
- Each file with 10 GridFTP streams
- 300 streams proved to be too much for some endpoints
  - PIC and RAL bandwidth stalled with 30 files
  - 10 files gave good throughput

Pre 19/10 many problems with Castor2/FTS interaction

- Files not staged cause FTS transfers to timeout/fail
  - Current not possible to transfer files from tape directly with FTS
  - Pre-staged files to disk - ~50k files for transfer (~75k in total: 10 TB)
- CASTOR2 too many problems to list ...
  - Reliability of service increased markedly when ORACLE server machine upgraded



## Problems...

### srm\_advisory\_delete

- Inconsistent behaviour of SRM depending on "backend" implementation
  - Not well - defined functionality in SRM v1.1
- Not possible to physically delete files in consistent way on the Grid at the moment
  - dCache can "advisory delete" and re-write - can't overwrite until an "advisory delete"
  - CASTOR can simply overwrite !

### FTS failure problems

- Partial transfer can't re-transfer after failure
  - FTS failed to issue an "advisory delete" after a failed transfer
  - Can't re-schedule transfer to dCache sites until an "advisory delete" issued manually

## Problems...

### LFC registration/query

- This is currently limiting factor in our system
  - Moving to using "sessions" - remove authentication overhead for each operation
    - Under evaluation
  - (another approach read-only insecure front-end for query operations)

# The ATLAS Experiment

# ATLAS SC3 goals



- Exercise ATLAS data flow
- Integration of data flow with the ATLAS Production System
- Tier-0 exercise
- Completion of a “Distributed Production” exercise
  - Has been delayed

**Following slides on Tier0 dataflow exercise which is running now!**

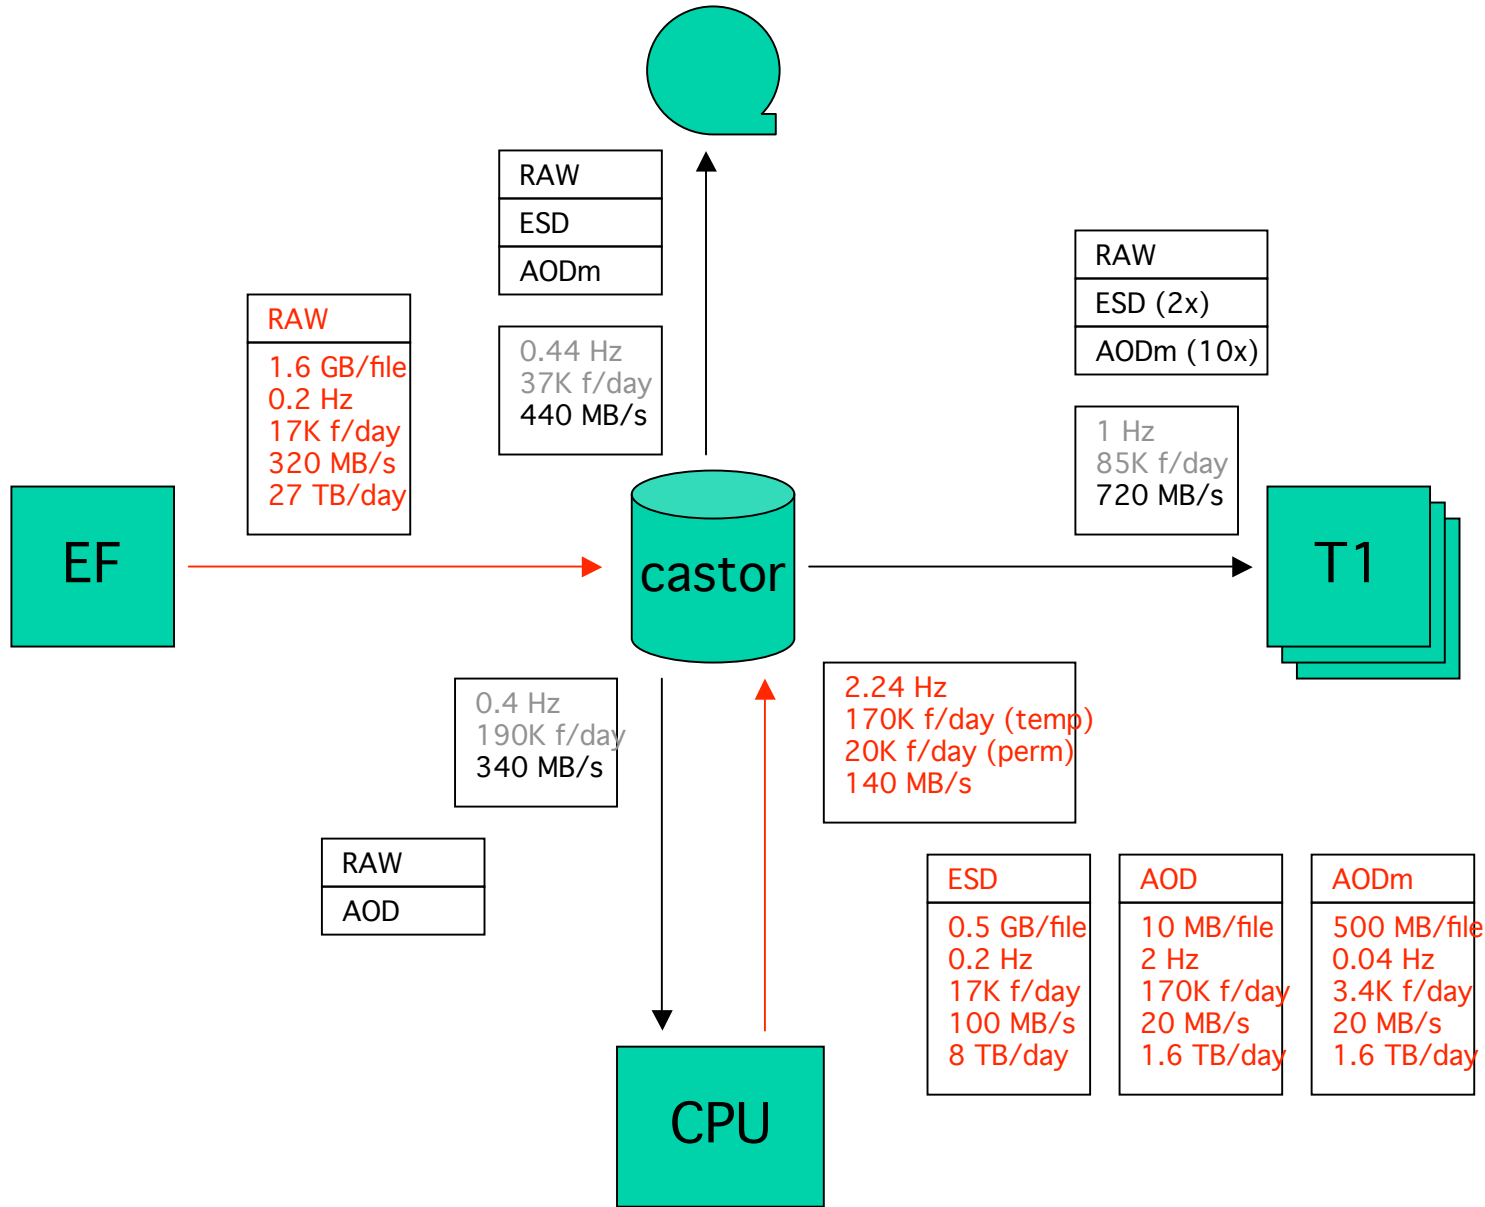
- More information:
  - <https://uimon.cern.ch/twiki/bin/view/Atlas/DDMSc3>



# ATLAS-SC3 Tier0

- Quasi-RAW data generated at CERN and reconstruction jobs run at CERN
  - No data transferred from the pit to the computer centre
- “Raw data” and the reconstructed ESD and AOD data are replicated to Tier 1 sites using agents on the VO Boxes at each site.
- Exercising use of CERN infrastructure ...
  - Castor 2, LSF
- ... and the LCG Grid middleware ...
  - FTS, LFC, VO Boxes
- ... and expt Distributed Data Management (DDM) software

# ATLAS Tier-0



# ATLAS-SC3 Tier-0



- Main goal is a 10% exercise
  - Reconstruct “10%” of the number of events ATLAS will get in 2007 using “10%” of the full resources that will be needed at that time
- Tier-0
  - ~300 kSI2k
  - “EF” to CASTOR: 32 MB/s
  - Disk to tape: 44 MB/s (32 for raw and 12 for ESD+AOD)
  - Disk to WN: 34 MB/s
  - T0 to T1: 72 MB/s
  - 3.8 TB to “tape” per day
- Tier-1 (in average):
  - ~1000 files per day
  - 0.6 TB per day
  - At a rate of ~7.2 MB/s

# SC3 pre-production testing

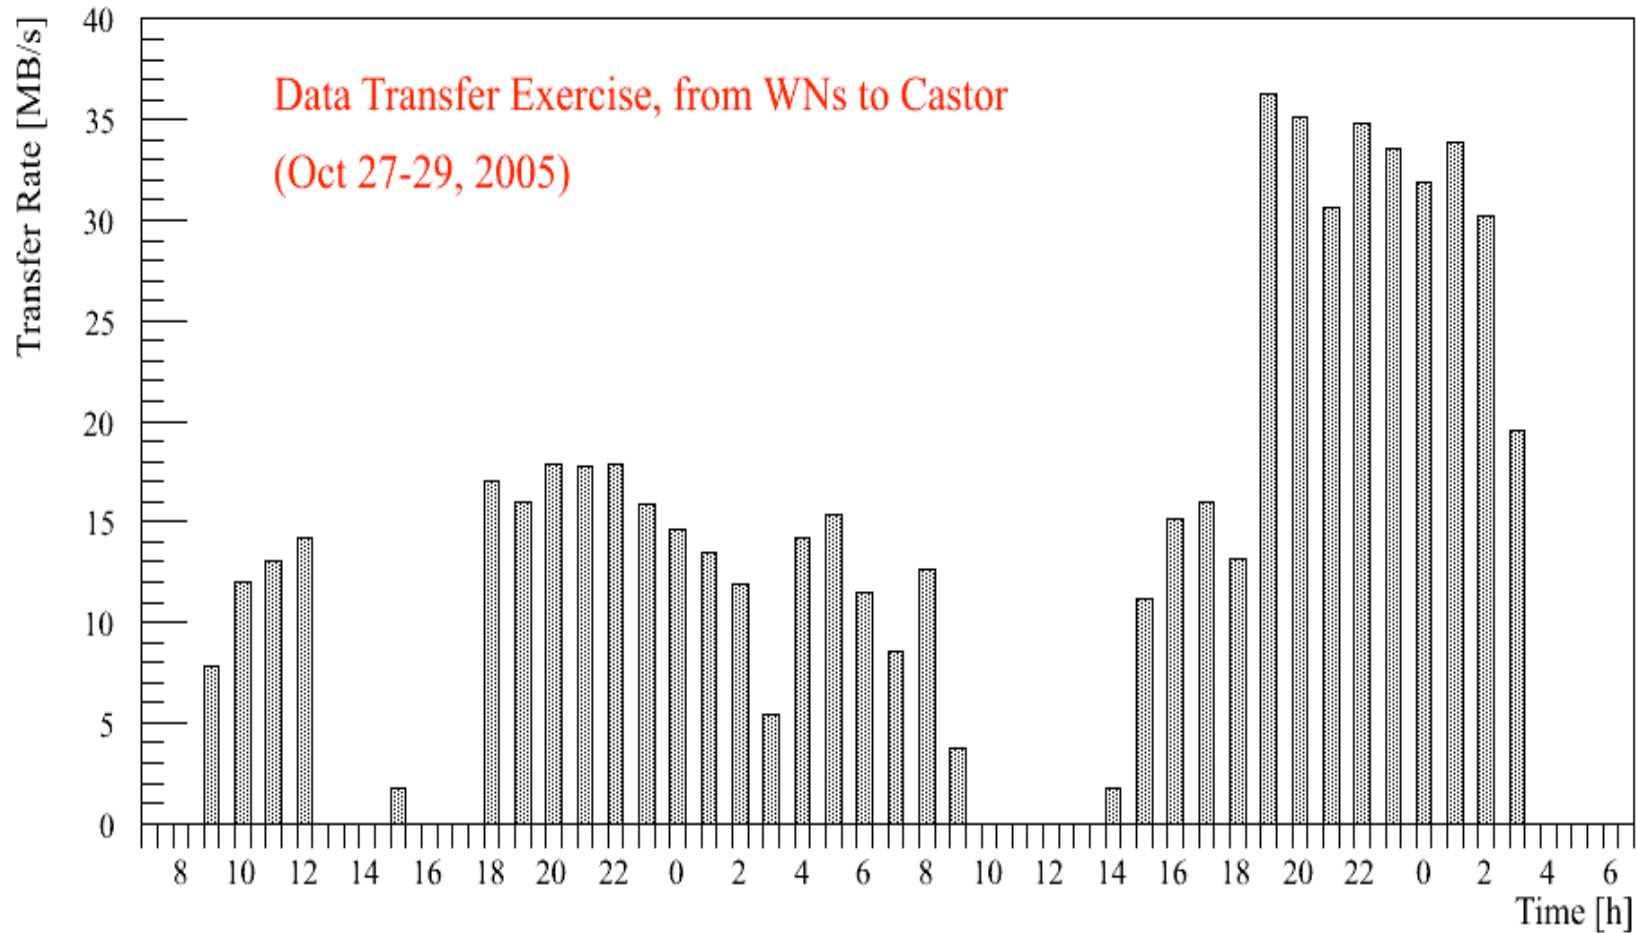


- For 2/3 weeks up to 1st November tested the functionality of SC3 services integrated with ATLAS DDM (FTS, LFC etc..)
- Very useful to have this testing phase with pilot services since there were many problems...
  - On expt side, good test of deployment mechanism - fixed bugs & optimised
  - on sites - mainly trivial like SRM paths etc.
  - A point on mailing lists...
    - There are too many
    - For official problem reporting: [lcg-sc.support@cern.ch](mailto:lcg-sc.support@cern.ch) - sometimes ticketing system doesn't work and response is slow
- Better coordination needed when deploying components to avoid conflicts (for example LCG/POOL/Castor)



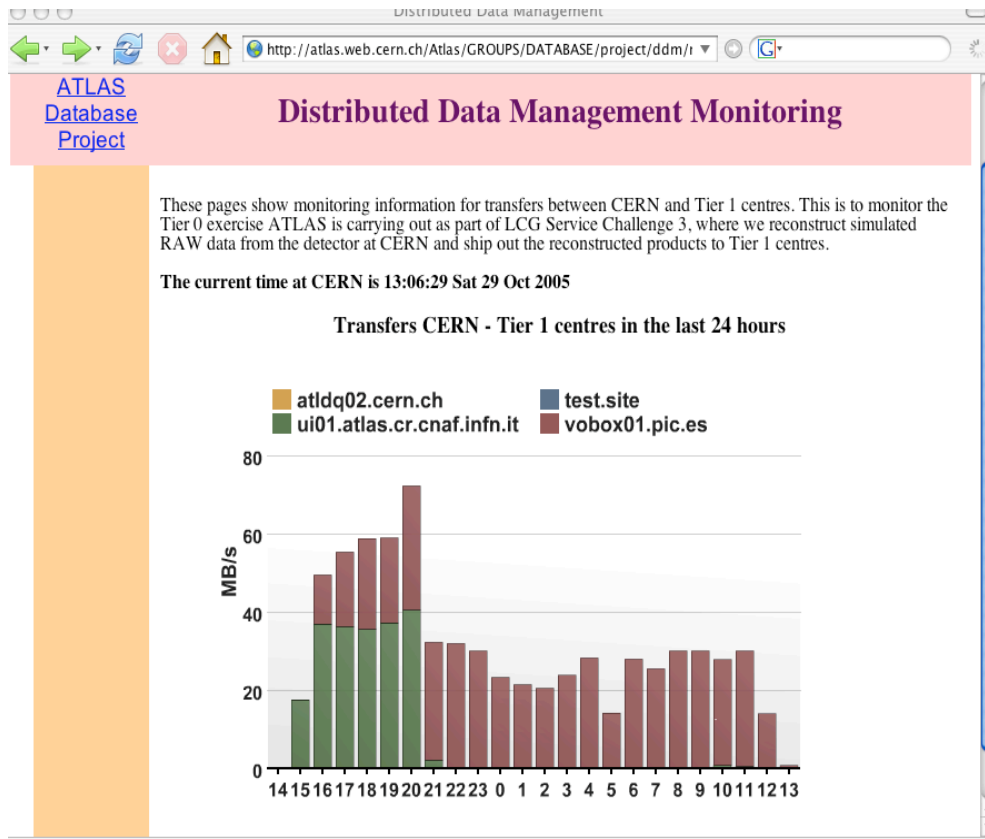


# Data transfer





# 24h before 4 day intervention 29/10 - 1/11

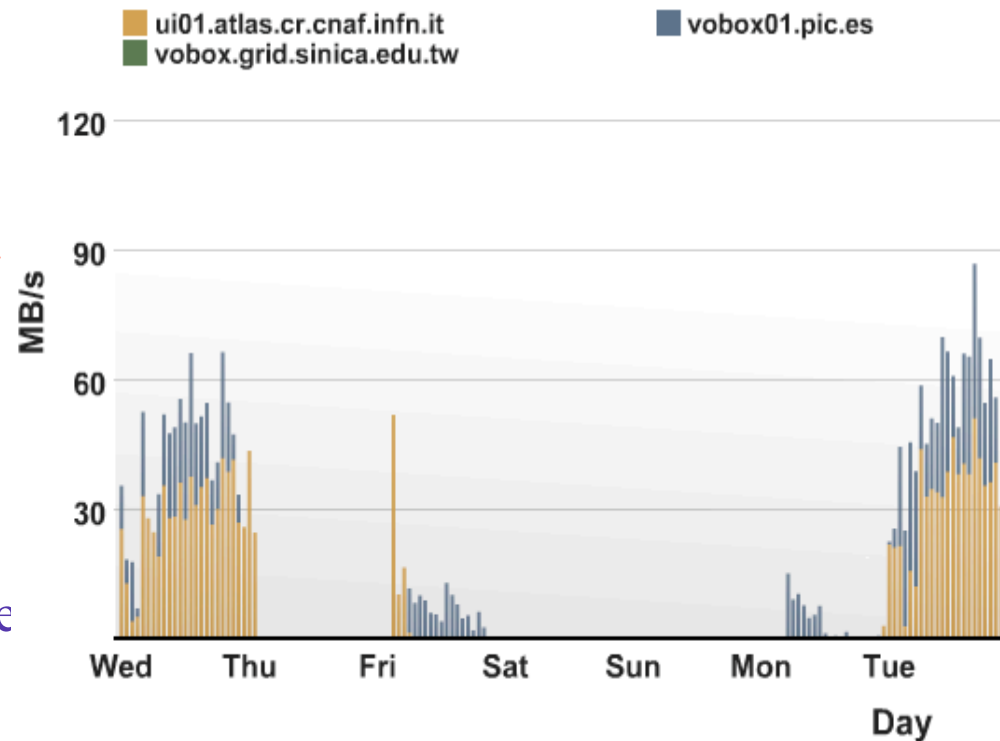


We achieved quite good rate in the testing phase (sustained 20-30 MB/s to one site (PIC))



# SC3 experience in 'production' phase

- Started on Wed 2nd Nov - ran smoothly for ~24h (above bandwidth target) until... problems occurred with all 3 sites simultaneously
  - CERN: power cut and network problems which then caused castor namespace problem
  - PIC: Tape library problem meant FTS channel switched off
  - CNAF: LFC client upgraded and not working properly
- It took about 1 day to solve all these problems
- No jobs running during the last weekend.





# Data Distribution

- Use a generated “dataset”
  - Contains 6035 files (3 TB) and we tried to replicate it to CNAF and PIC
- PIC: 3600 files copied and registered
  - 2195 ‘failed replication’ after 5 retries by us x 3 FTS retries
    - Problem under investigation
  - 205 ‘assigned’ - still waiting to be copied
  - 31 ‘validation failed’ since SE is down
  - 4 ‘no replicas found’ LFC connection error
- CNAF: 5932 files copied and registered
  - 89 ‘failed replication’
  - 14 ‘no replicas found’

# General view of SC3



- When everything is running smoothly ATLAS get good results
- The middleware (FTS, LFC) is stable but the sites' infrastructure is still very unreliable
  - ATLAS DDM software dependencies can also cause problems when sites upgrade middleware
- good response from LCG and sites when there are problems (proviso earlier email list comment)
  - But the sites are not running 24/7 support
  - Means a problem discovered at 6pm on Friday may not be answered until 9am on Monday so we lose 2 1/2 days production
- Good cooperation with CERN-IT Castor and LSF teams.
- not managed to exhaust anything production s/w; LCG m/w)
- Still far from concluding the exercise and not running stably in any way - cause for concern
- Exercise will continue adding new sites.

## General Summary of SC3 experiences

Reliability seems to be the major issue:

- CASTOR2 - still ironing out problems, but big improvements in service
- Coordination issues
- Problems with sites and networks
  - MSS, security, network, services...

FTS:

- For well-defined site/channels performs well after tuning
- Timeout problems dealing with accessing data from MSS
- Clean-up problems after transfer failures
- Ability for a centralised service for 3rd party transfers
- Plenty to discuss at next week's workshop

SRM:

- Limitations/ambiguity (already flagged) in functionality