Expt SC3 Status

Nick Brook

In chronological order:

ALICE

CMS

LHCb

ATLAS

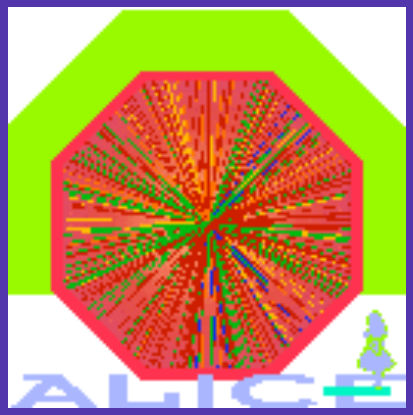

Alice Physics Data Challenge '05 - goals

- 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

Alice Physics Data Challenge '05 - goals

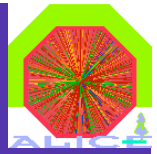

- Data production:
  - List of physics signals defined by the ALICE Physics Working Groups
  - Data used for detector and physics studies
  - Approximately 500K Pb+Pb events with different physics content, 1M p+p events, 80TB production data and few TB user generated data
  - Structure – divided in three 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 the data can pushed 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
- 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, FZK (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

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
- Good collaboration with IT-GD/FIO groups with the installation and operations …
- … 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

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

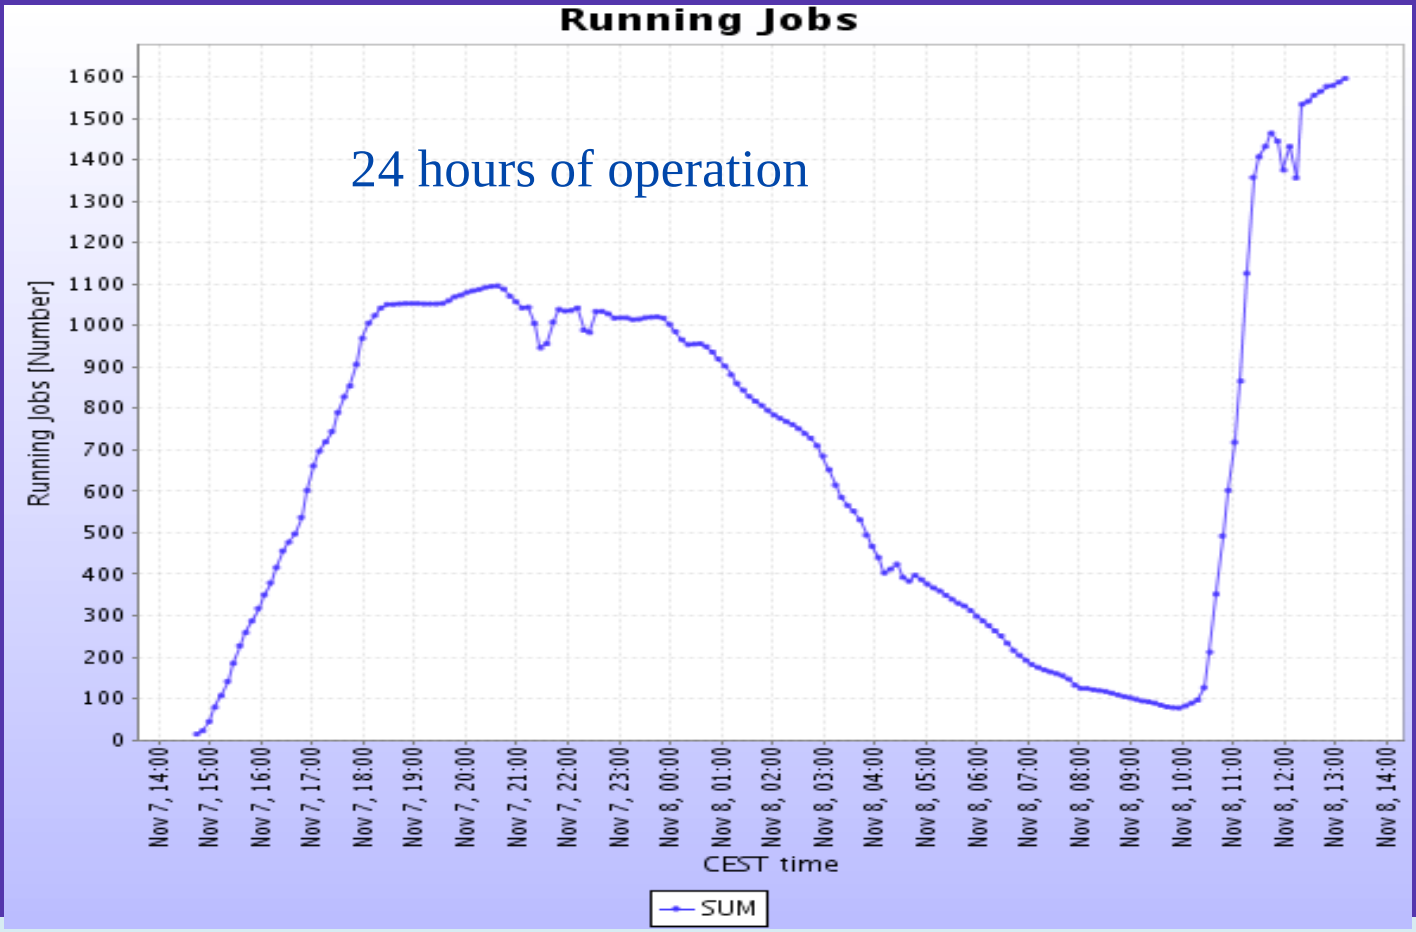

Physics Data Challenge

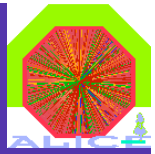

- T1s provide the required services and the announced resources
- Good cooperation with LCG and all the major T1s management
- A 10 days statistics (starting phase) ~ 8000 jobs done
  - FZK: 32%
  - CERN: 21%
  - CCIN2P3: 16%
  - CNAF: 12% (started slightly later)
  - NIKHEF: 0% (problems keeping up with s/w updates)
  - RAL: 0% (VO-box in preparation)
  - NDGF: 0% (working on interface)
  - The remainder provided by T2s (Italy, Germany, Russia, France, Czech Republic, South Korea, Romania, Poland, India) entering progressively the exercise

ALICE plans:

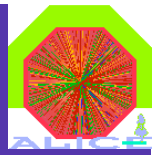

- File replication with FTS:
  - FTS endpoints tested at all ALICE SC3 sites
  - Start data migration in about 10 days, initially T0->T1
  - Test, if possible, migration Tx->Ty
- Re-processing of data with calibration at T0/T1:
  - AliRoot framework ready, currently calibration and alignment algorithms 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


CMS

CMS SC3 Goals and Operations

- Integration test of data transfer and data serving infrastructure
  - Computing Integration Test exercising the bulk data processing portion of the CMS computing model under realistic conditions
    - Test end-to-end system of CMS-specific and LCG services
    - Focused validation of data storage, transfer and serving infrastructure, plus required workload components: job submission, resource broker etc.
- Test complexity built up in three major steps over 2005
  - Throughput phase (June): high-throughput storage+transfer system test
  - Service I (Sep-Oct): Concurrent transfer+grid jobs to read/write data
  - Service II (Nov-Dec): Concurrent data flow scenario à la C-TDR
    - Tier-0 (simulated) raw/reco data —> Tier-1
    - Tier-1 skim production —> Tier-2s for skim analysis;
    - Tier-2 MC production data —>Tier-1
- Involve a significant number of Tier 1 and Tier 2 sites
    - CERN + all 7 CMS Tier 1s: ASGC, CNAF, FNAL, FZK, IN2P3, PIC, RAL
    - 13 Tier 2s: DESY (FZK); Bari, Legnaro (CNAF); CIEMAT+IFCA (PIC); NCU (ASGC); Imperial (RAL); Caltech, Florida, Nebraska, Purdue, UCSD, Wisconsin (FNAL); [Plus joining: SINP+ITEP]

SC3 Operations

- CMS central responsibilities
  - Data transfers entirely managed through PhEDEx by central transfer management database operated by PhEDEx operations
    - Using underlying grid protocols srmcp, globus-url-copy and FTS
    - Placing files through SRM on site storage based on Castor, dCache, DPM
  - CMS analysis jobs submitted by job robot through CMS CRAB system
    - Using LCG RB (gdrb06.cern.ch) and OSG Condor-G interfaces
  - monitoring info centrally collected using MonaLisa and CMS Dashboard
    - Fed from RGMA, MonALISA and site monitoring infrastructure
- Site responsibilities (by CMS people at or “near” site)
  - ensuring site mass storage and mass storage interfaces are functional, grid interfaces are responding, and data publishing steps are succeeding
    - Data publishing, discovery: RefDB, PubDB, ValidationTools
    - Site local file catalogues: POOL XML, POOL MySQL
  - A lot of infrastructure tools are provided to the sites, but having the whole chain hang together requires perseverance

Service Phase I Results

- SC3 Phase I: transferred data and & processing jobs
  - Total volume transferred in Sep-Oct: 145 TB
    - Roughly as much as CMS has transferred in the last 12 months
    - Details on data transfer volumes and succes rates in tables below
  - ran several thousands of jobs

T1 Site	Volume	Quality	Hours	Rate
ASCC	20.9 TB	5%	604	10.1 MB/s
CNAF	9.6 TB	7%	514	5.4 MB/s
FNAL	47.0 TB	39%	1060	12.9 MB/s
FZK	9.8 TB	23%	646	4.4 MB/s
IN2P3	1.2 TB	1%	309	1.1 MB/s
PIC	3.1 TB	32%	120	7.5 MB/s
RAL	6.8 TB	14%	425	4.7 MB/s

T2 Site	Volume	Quality	Hours	Rate
Bari	4.0 TB	71%	227	5.1 MB/s
Caltech	2.2 TB	0%	666	1.0 MB/s
DESY	3.5 TB	1%	378	2.7 MB/s
Florida	3.0 TB	24%	204	4.3 MB/s
Legnaro	3.6 TB	90%	82	12.8 MB/s
NCU	1.9 TB	2%	331	1.7 MB/s
Nebraska	13.8 TB	4%	682	5.9 MB/s
Purdue	6.5 TB	12%	223	8.5 MB/s
Spain	1.4 TB	59%	48	8.5 MB/s
UCSD	1.9 TB	83%	104	5.3 MB/s
Wisconsin	1.9 TB	0%	723	0.8 MB/s

Quality = Successful transfers vs. those started
Hours = Number of hours with successful transfers
Rate = Volume / Hours

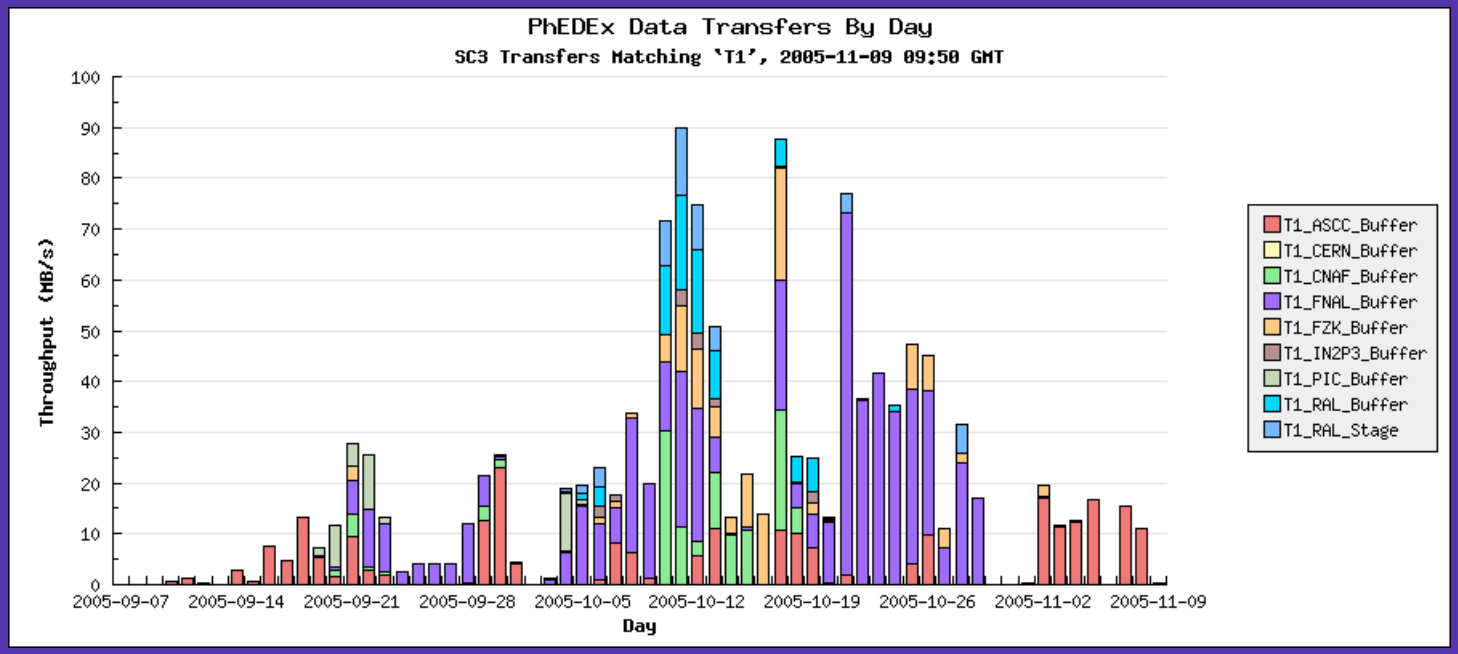

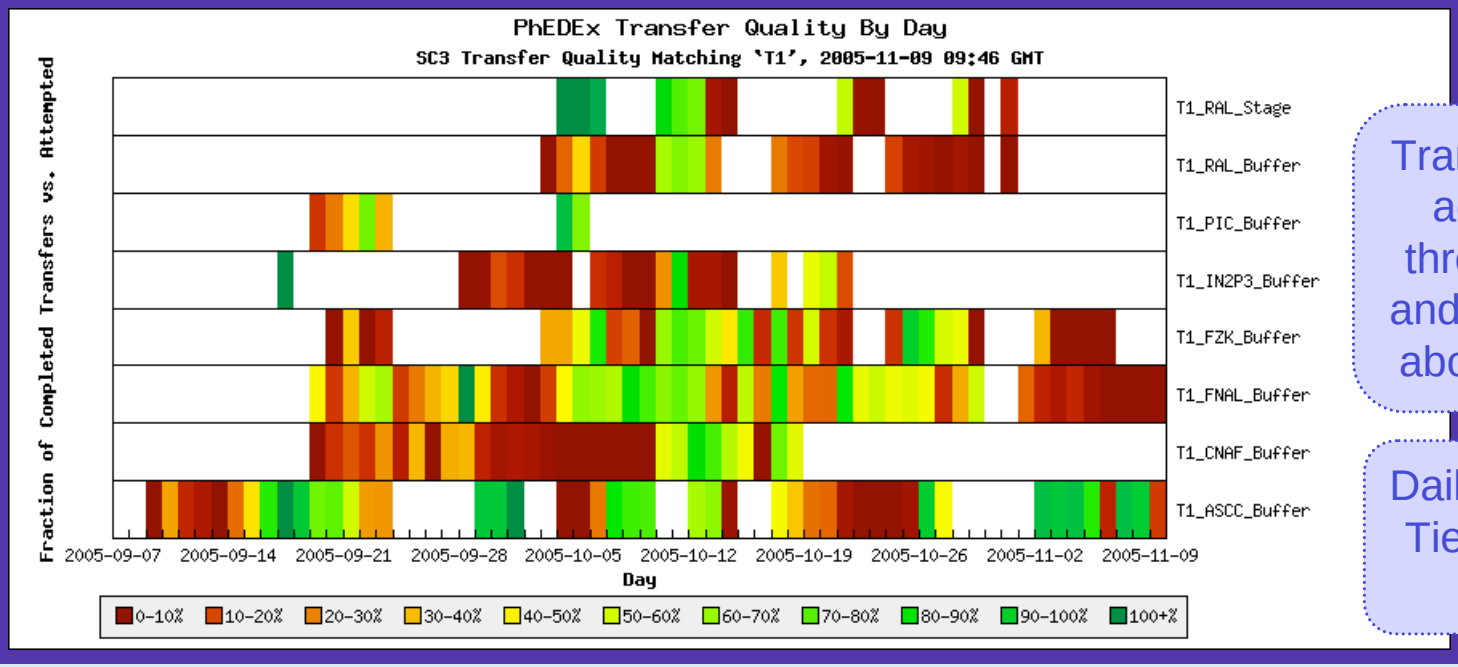

Tier-1 WAN Transfers

Transfers were balanced across several sites, throughput was healthy and error rate modest for about **5 days out of 50.**

Daily aggregate rate from Tier 0 to Tier 1s peaked at **90 MB/s.**

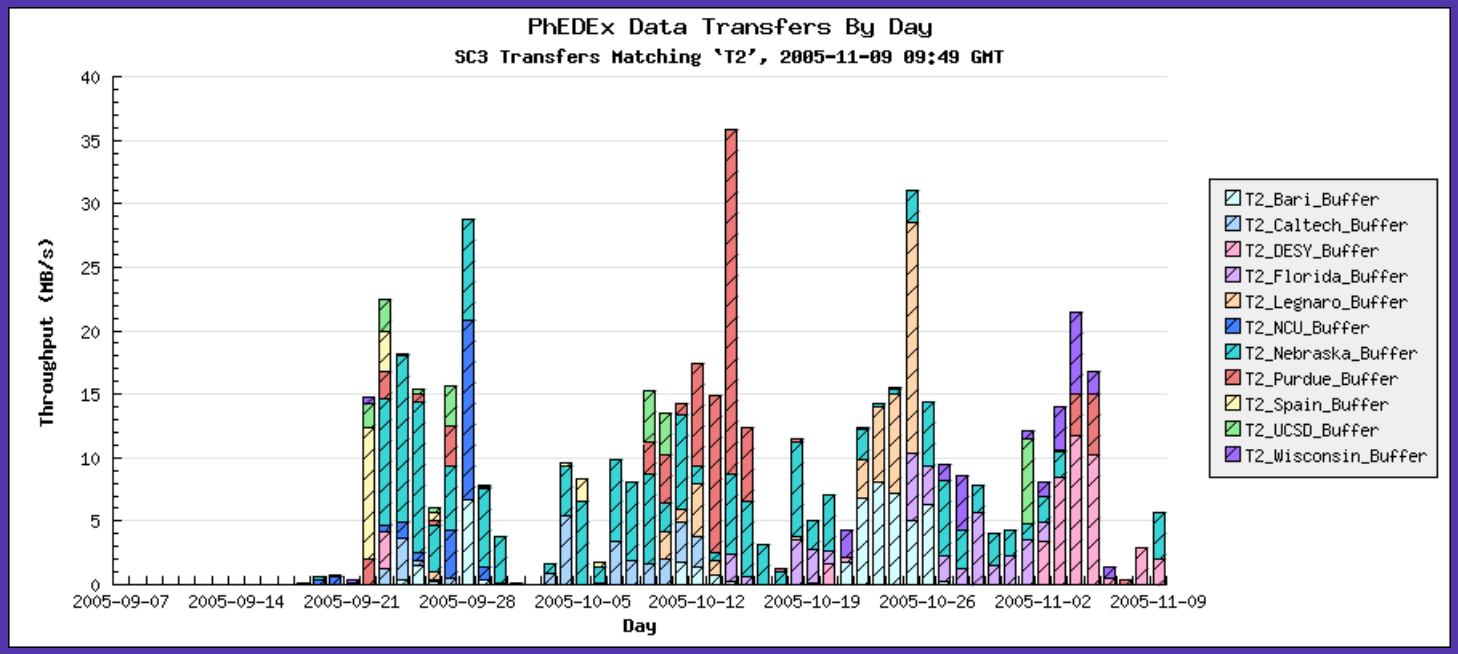

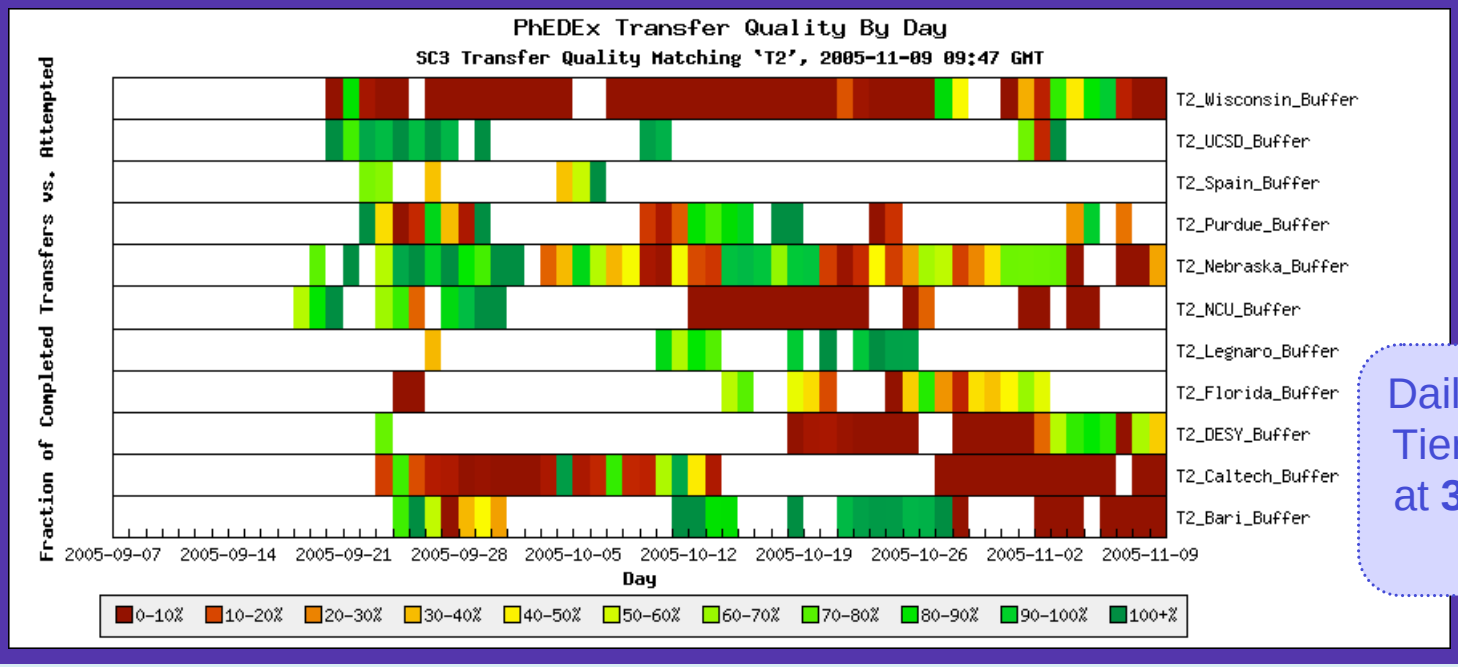

Tier-2 WAN Transfers

Daily aggregate rate from Tier 1s to Tier 2s peaked at **35 MB/s**, typically well below **20 MB/s**.

Summary of Experiences

- Months of intense debugging is beginning to bear fruit
  - Promising results and impressive effort by numerous sites, but…
  - debugging and shaking out components overwhelmed end-to-end goals
    - Many services were not sufficiently tested before start of challenge period
    - De-scoped to debugging pieces that did not work as expected.
- Lessons learned and principal concerns
  - Castor-2: Innumerable problems, we now hope to run more smoothly
  - SRM: Less standard than anticipated, lacking tuning at Castor/SRM sites
  - LFC: integration work was done for use as CMS/POOL file catalog
  - DPM: RFIO incompatibilities make CMS applications fail to access files
  - FTS: Integration ongoing, move to FTS 1.4
  - CMS data publishing: Difficult to configure and very difficult to operate
    - Looking forward to improvements with new system
  - CMS software releases: Improve release/distribution process,validation

SC3 Plans for Phase2

- Hope to complete limited-scope integration test until end of 2005
  - Fall back to subset of highest priority objectives
    - demonstrate stable transfers from Tier-0 to Tier-1 centers and understand and reduce the failure rates
    - demonstrate reasonable throughput out of local mass storage under prototypical analysis applications
  - Begin in second half of November, December as reserve
    - Other tests will need to be revisited in CMS Integration Program outside SC3
- SC3 has been costly
  - Substantial efforts by few individuals
  - CMS development program reduced to allocate effort to debug SC3
  - Several Tier 2s unable to secure attention from Tier 1s swamped by SC3
  - Re-evaluation of CMS integration plans to address issues & delays
  - As prepare for operations - need to arrive at a state where the challenge types of activities are becoming more mundane
- Service Challenges enormously important to establish WLCG service!


LHCb

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 central machine at CERN

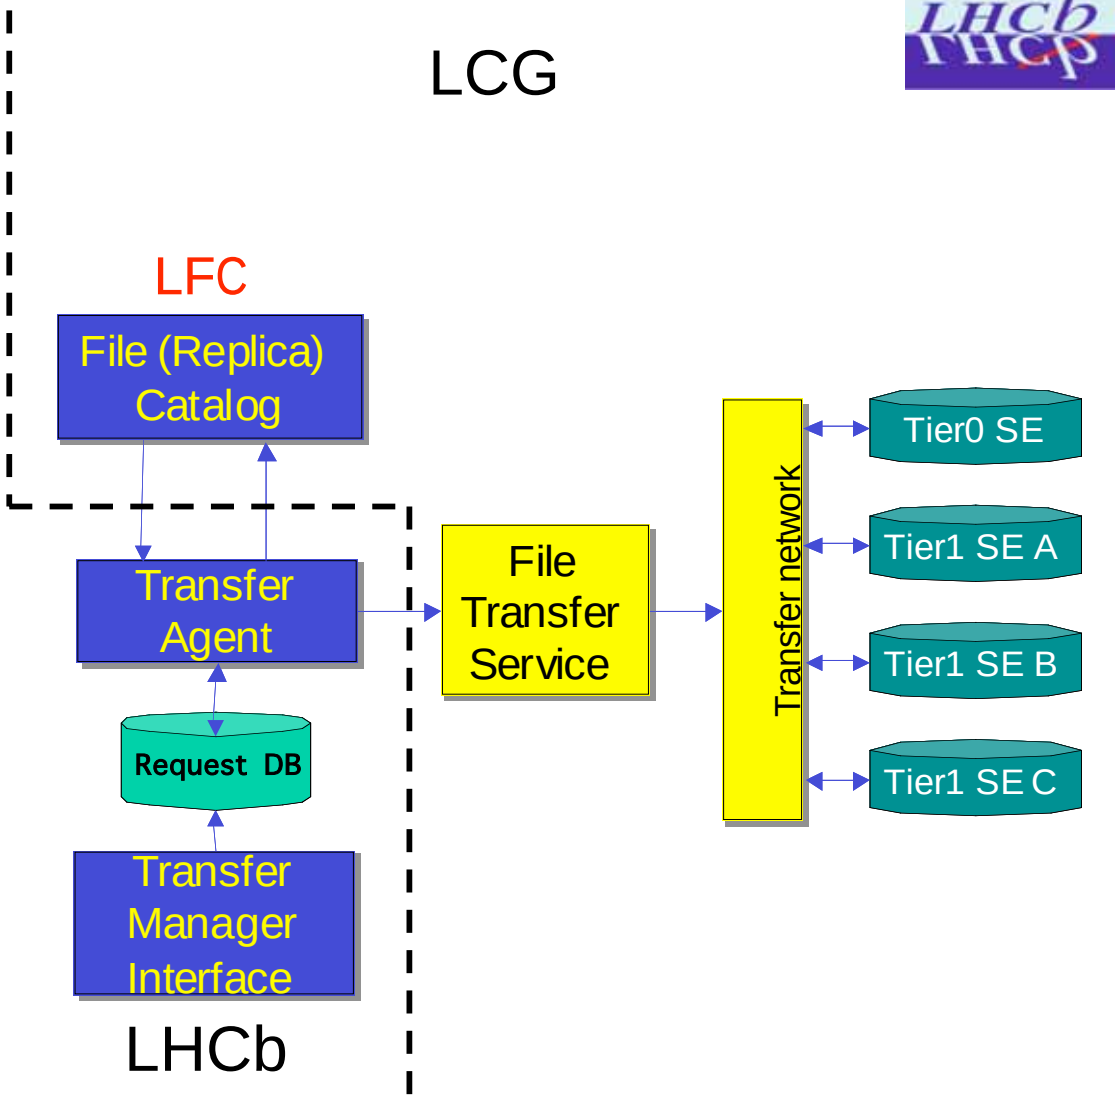

DIRAC transfer agent

- 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

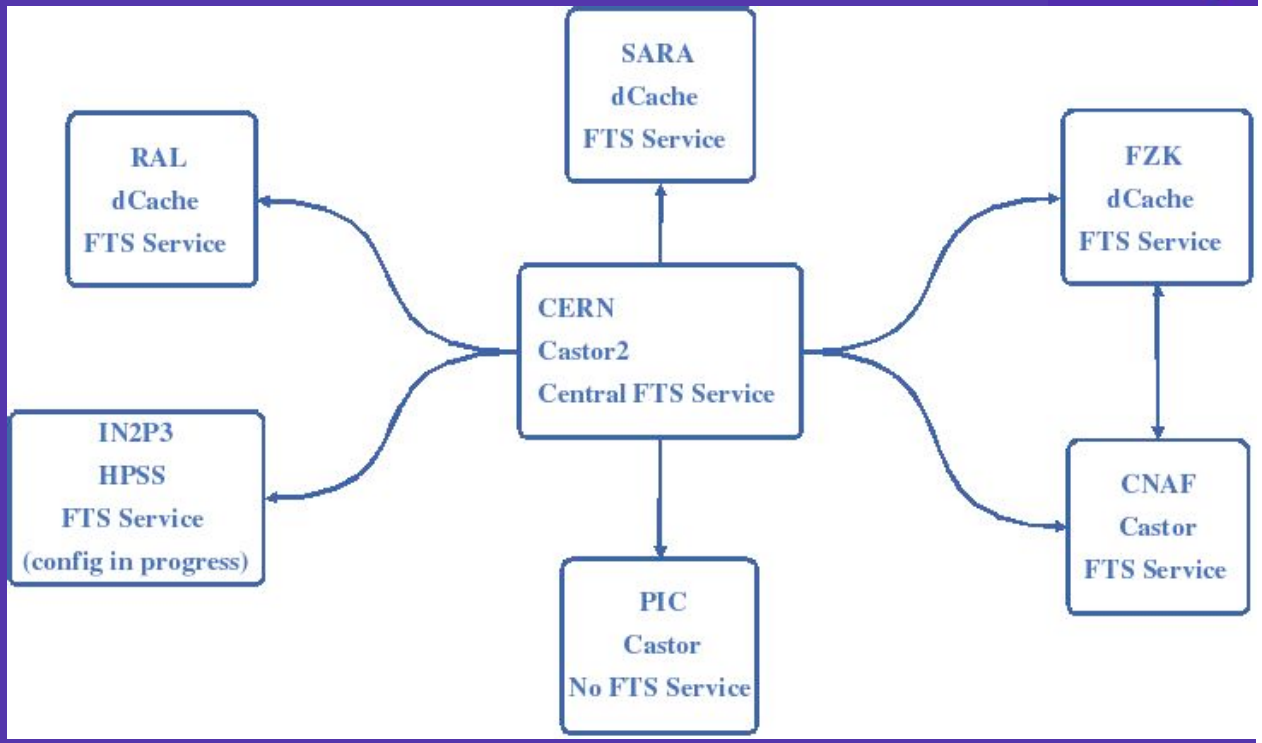

FTS central service for managing T1-T1 matrix ??

Overview of SC3 activity

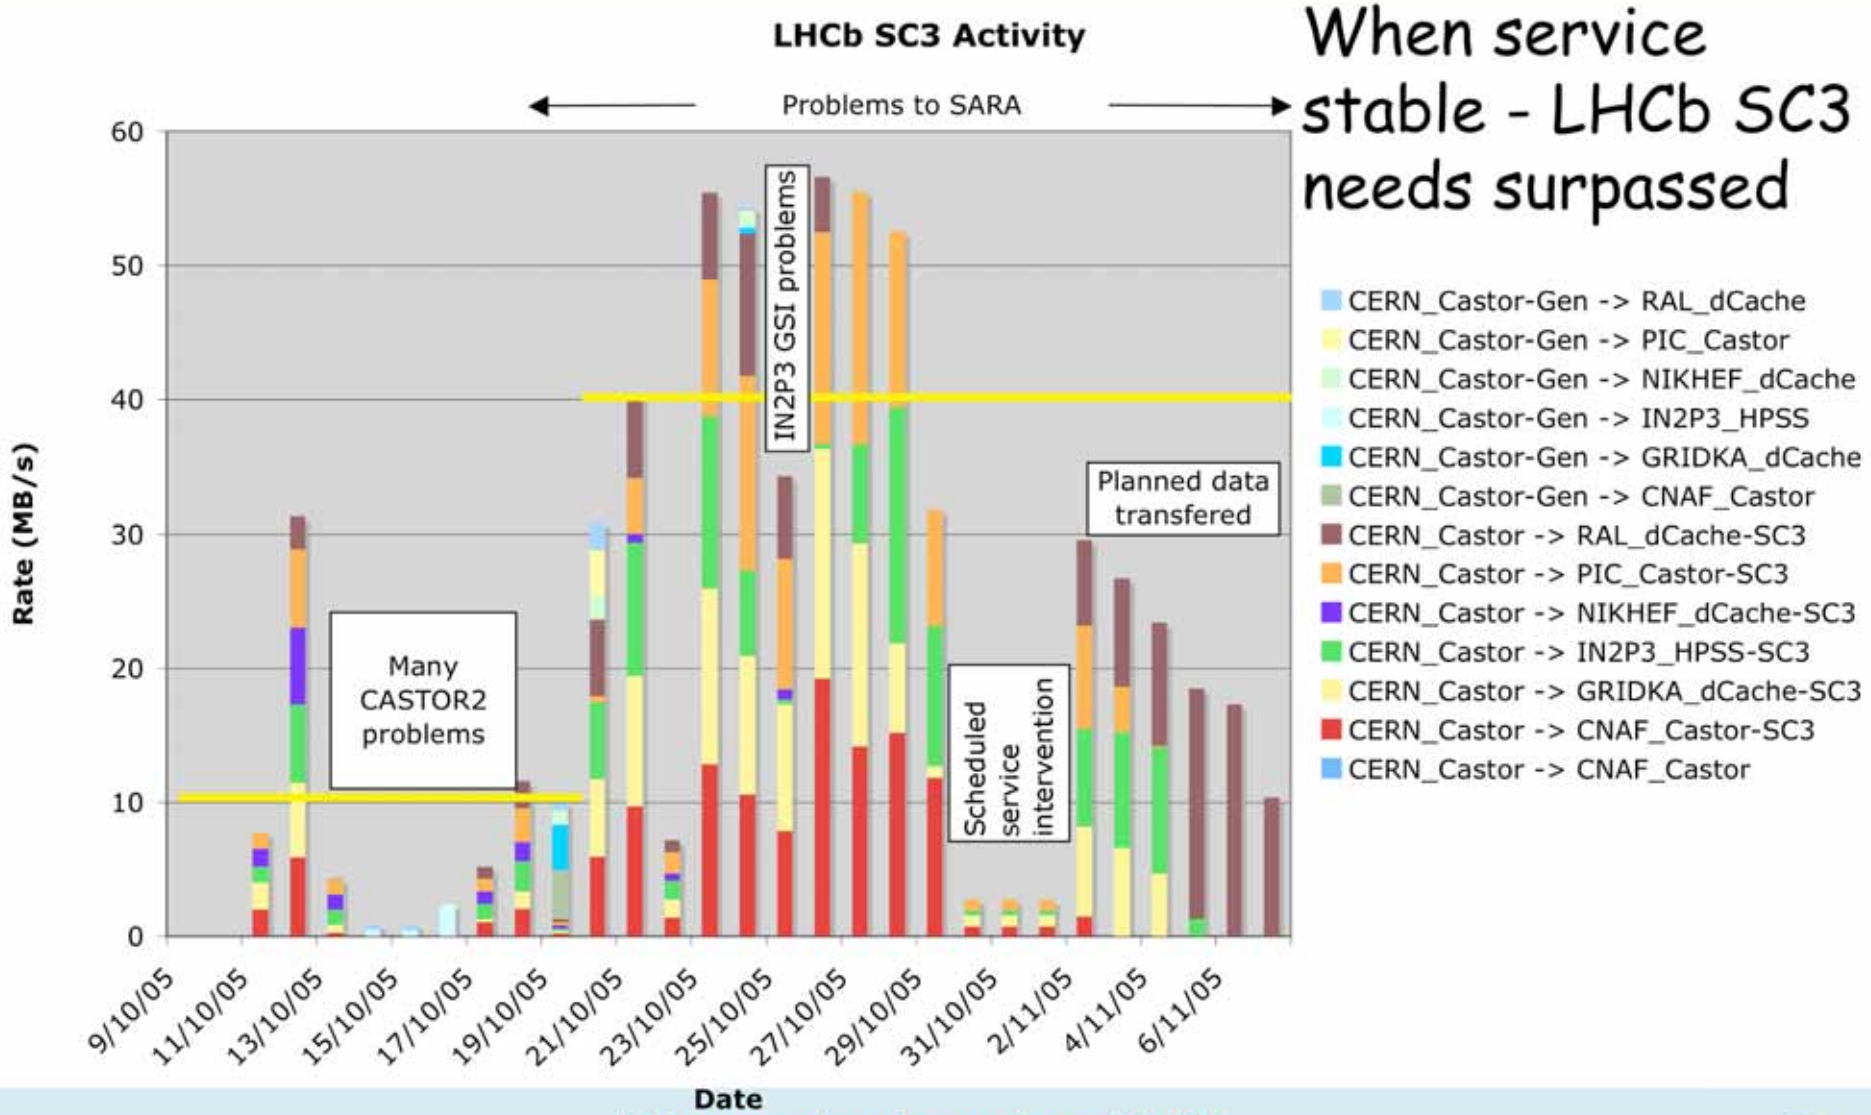

Experiences…

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

Early October 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

Experiences…

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

Experiences…

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)

Good interaction with FTS, LFC, CASTOR-2 teams

Sites very supportive

The ATLAS Experiment 

ATLAS SC3 goals

- Exercise ATLAS data flow
- Integration of data flow with the ATLAS Production System
- Tier-0 exercise
- "Distributed Production" exercise
  - Will come afterwards

Concentrate on Tier0 dataflow exercise which is running now!

- More information:
  - https://uimon.cern.ch/twiki/bin/view/Atlas/DDMSc3
- Real-time monitoring of data transfers on:
  - http://atlas-ddm-monitoring.web.cern.ch/atlas-ddm-monitoring/
- ATLAS Distributed Data Management:
  - https://uimon.cern.ch/twiki/bin/view/Atlas/DDM

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 software
  - Production System: new Supervisor (Eowyn)
  - Tier0 Management System (TOM)
  - Raw Data generator (Jerry)
  - Distributed Data Management (DDM) software (DQ2)

Dataflow 2007

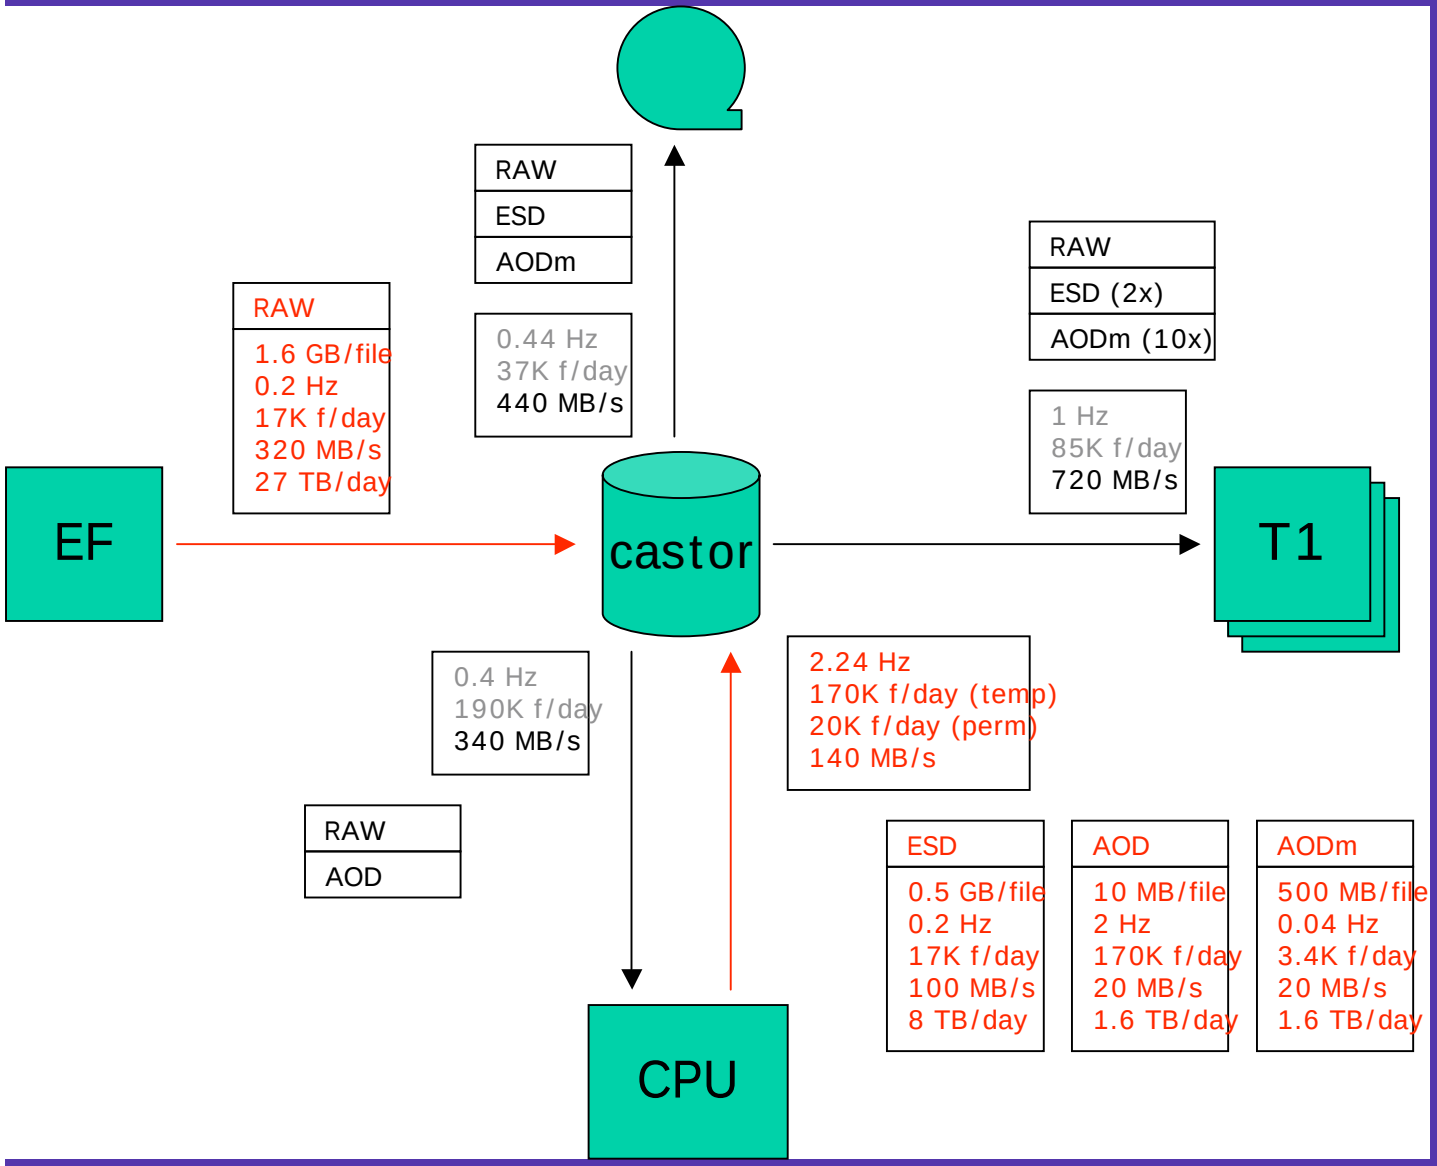

SC3 10% challenge of 2007 rates

Snapshot of Activity

24h period on Wednesday, November 9

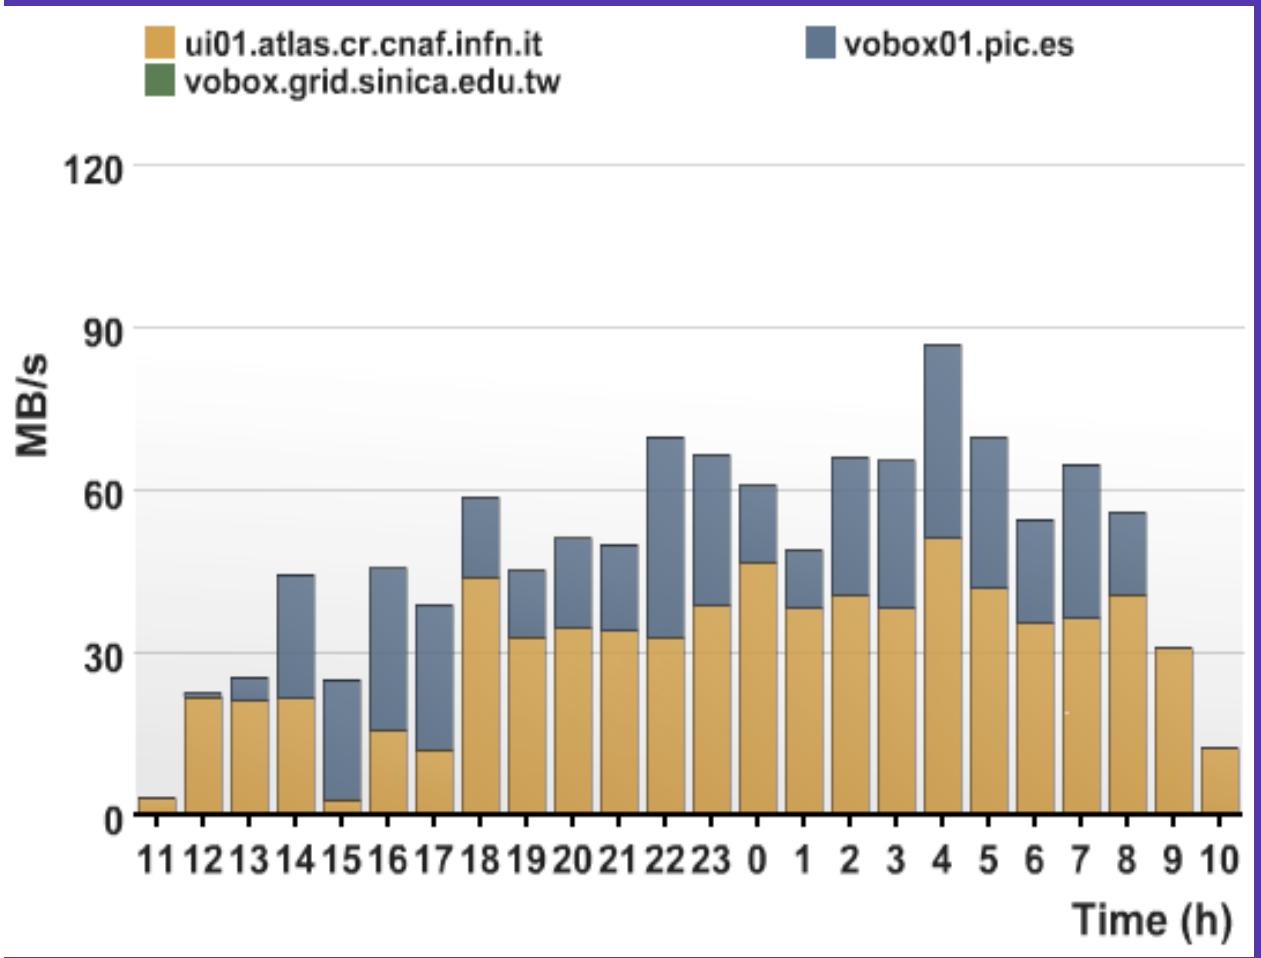

achieved quite good rate (sustaining 20-30 MB/s to sites)

Pretty much accomplished the expected **rate for real data taking** for CNAF and 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 weekend (5/6th November)

Transfers "CERN to Tier1 centres"
Average throughput per hour
November 2-7

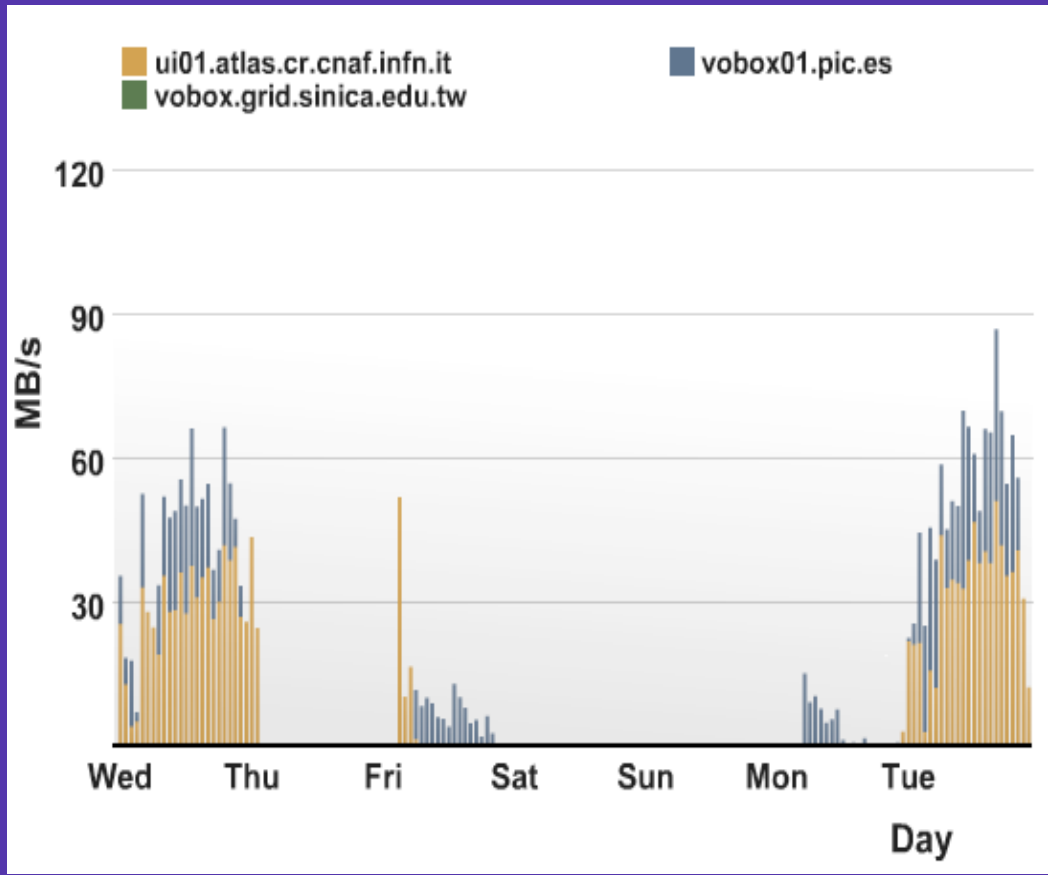

SC3 experience in 'production' phase

More Tier1s are now joining the exercise…

Now:

- **BNL**
- **CNAF**
- **SARA**
- **PIC**
- **ASGC**

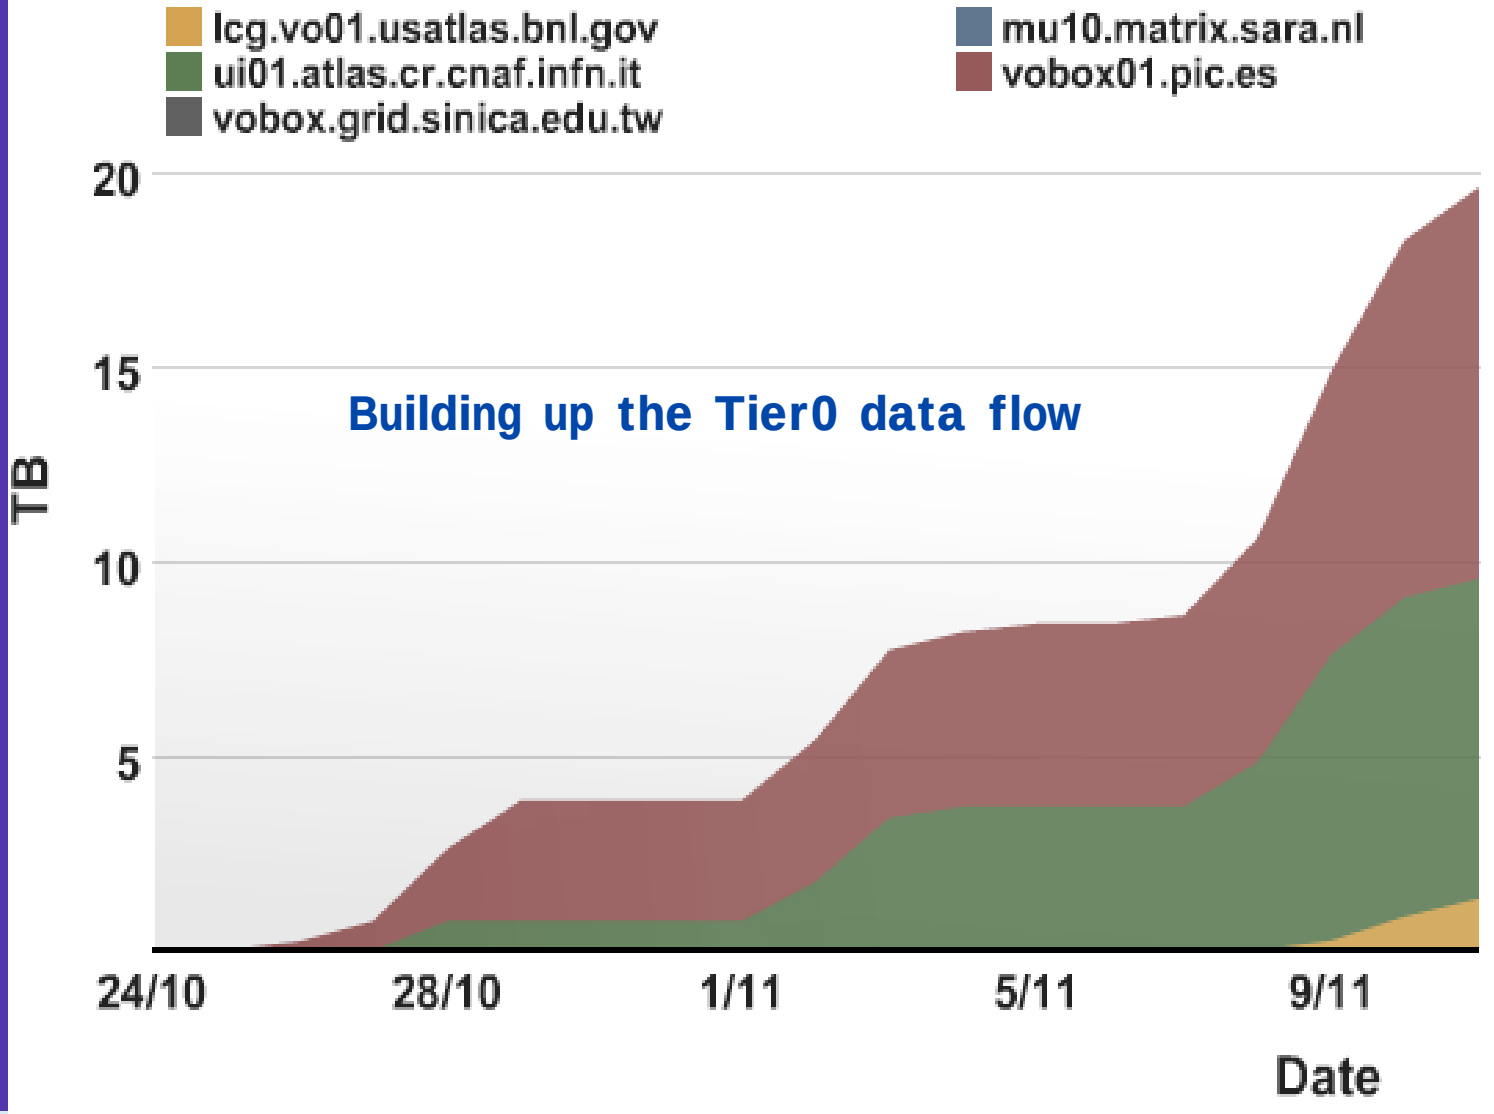

General view of SC3

- When everything is running smoothly - good results
- The middleware (FTS, LFC) is stable but instability in the sites' infrastructure
  - ATLAS DDM software dependencies can also cause problems when sites upgrade middleware
- good response from LCG and sites when there are problems 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
- Exercise will continue adding new sites

General Summary of SC3 experiences

Extremely useful for shaking down sites, experiment systems & WLCG

- Many new components used for the 1st time in anger
- Need for additional functionality in services
  - FTS, LFC, SRM, …

Reliability seems to be the major issue

- CASTOR2 - still ironing out problems, but big improvements
- 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

SRM:

- Limitations/ambiguity (already flagged) in functionality