

SC3 - Experiments' Experiences

Nick Brook

In chronological order:

ALICE


CMS

LHCb

ATLAS




General running statistics




- Event sample (last two months running)
 - 22,500 jobs completed (Pb+Pb and p+p)
 - Average duration 8 hours, 67,500 cycles of jobs
 - Total CPU work: 540 kSi2k hours
 - Total output: 20 TB (90% CASTOR2, 10% Site SE's)
- Centres participation (22 total)
 - 4 T1's: CERN, CNAF, GridKa, CCIN2P3
 - 18 T2's: Bari (I), Clermont (FR), GSI (D), Houston (USA), ITEP (RUS), JINR (RUS), KNU (UKR), Muenster (D), NIHAM (RO), OSC (USA), PNPI (RUS), SPbSU (RUS), Prague (CZ), RMKI (HU), SARA (NL), Sejong (SK), Torino (I), UiB (NO)
- Jobs done per site
 - T1's: CERN 19%, CNAF 17%, GridKa 31%, CCIN2P3 22%
 - Very evenly distribution among the T1's
 - T2's: total of 11%
 - Good stability at: Prague, Torino, NIHAM, Muenster, GSI, OSC
 - Some under-utilization of T2 resources - more centres available, could not install the Grid software to use fully

Methods of operation

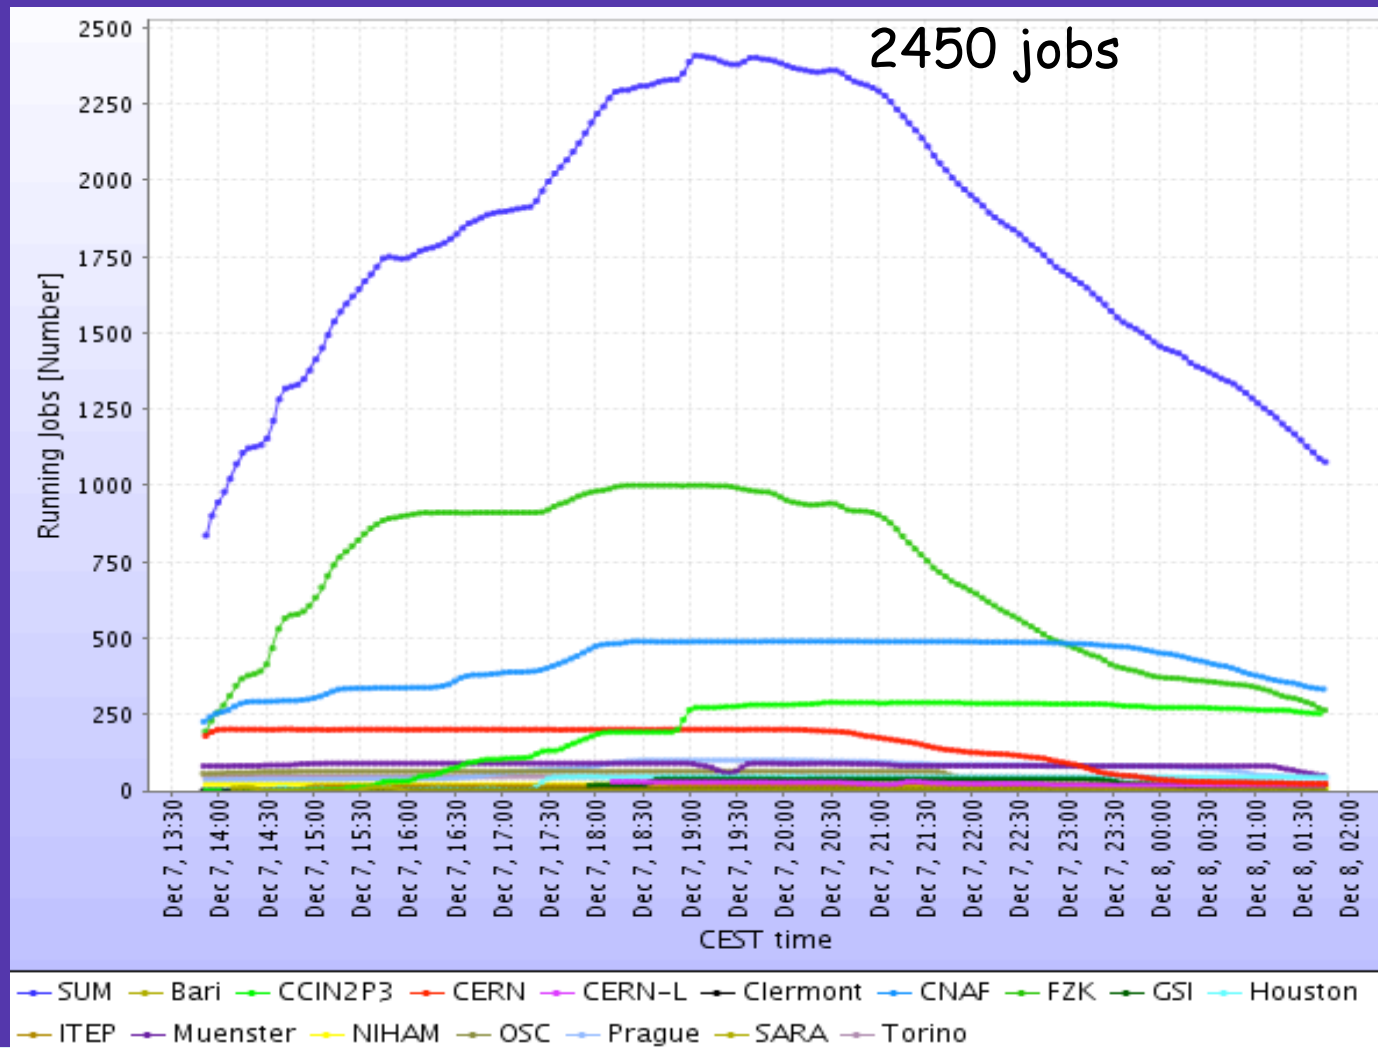


- Use LCG/EGEE SC3 baseline services:
 - Workload management
 - Reliable file transfer (FTS)
 - Local File Catalogue (LFC)
 - Storage (SRM), CASTOR2
- Run entirely on LCG resources:
 - Use the framework of VO-boxes provided at the sites


Status of production



Production job duration: 8 ½ hours on 1KSi2K CPU, output archive size: 1 GB (consists of 20 files)




Results




- VO-box behaviour

- No problems with services running, no interventions necessary
- Load profile on VO-boxes - in average proportional to the number of jobs running on the site, nothing special



ALICE VO-specific and LCG software



- Positive
 - Stable set of central services for production (user authentication, catalogue, single task queue, submission, retrieval of jobs)
 - Well established (and simple) installation methods for the ALICE specific VO-Box software
 - Good integration with the LCG VO-Box stack
 - Demonstrated scalability and robustness of the VO-Box model
 - Successful mass job submission through the LCG WMS
- Issues
 - Rapid updates of the MW problematic with inclusion of more computing centres on a stable basis
 - However all centres in the SC3 plan were kept up-to-date
 - Essentially due to a limited number of experts, currently focused on the software development
 - Not all services thoroughly tested, in particular LFC and FTS



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



Data Transfers

- ▶ Substantial effort on getting transfers to perform
 - * Total volume transferred in all of SC3 approximately 290 TB
 - * Total volume in the SC3 service phase 2: 140 TB
 - ◆ Roughly as much as in the 12 preceding months of production transfers!
 - ◆ 78 TB ended on tape in the period: 88% of Tier-1s–ASGC+DESY transfers
 - * Details on phase 2 data transfer rates and quality below

T1 Site	Volume	Quality	Hours	Rate
ASGC	19.7 TB	89%	323	17.8 MB/s
CNAF	7.8 TB	82%	124	18.3 MB/s
FNAL	23.6 TB	76%	402	17.1 MB/s
FZK	7.1 TB	35%	251	8.2 MB/s
IN2P3	19.4 TB	35%	276	20.5 MB/s
PIC	12.9 TB	62%	329	11.4 MB/s
RAL(*)	10.5 TB	18%	282	10.8 MB/s

T2 Site	Volume	Quality	Hours	Rate
Bari	3.4 TB	37%	179	5.6 MB/s
Caltech	1.1 TB	46%	178	1.8 MB/s
DESY	6.4 TB	59%	283	6.6 MB/s
Florida	1.5 TB	54%	22	20.4 MB/s
Legnaro	1.9 TB	27%	93	5.9 MB/s
NCU	2.9 TB	51%	82	10.2 MB/s
Nebraska	4.2 TB	5%	326	3.7 MB/s
Purdue	2.6 TB	43%	129	5.8 MB/s
Spain	2.1 TB	63%	96	6.5 MB/s
UCSD	5.8 TB	86%	172	9.8 MB/s
Wisconsin	1.9 TB	27%	257	2.2 MB/s

Quality = Successful transfers vs. those started

Hours = Number of hours with transfer activity

Rate = Volume / Hours

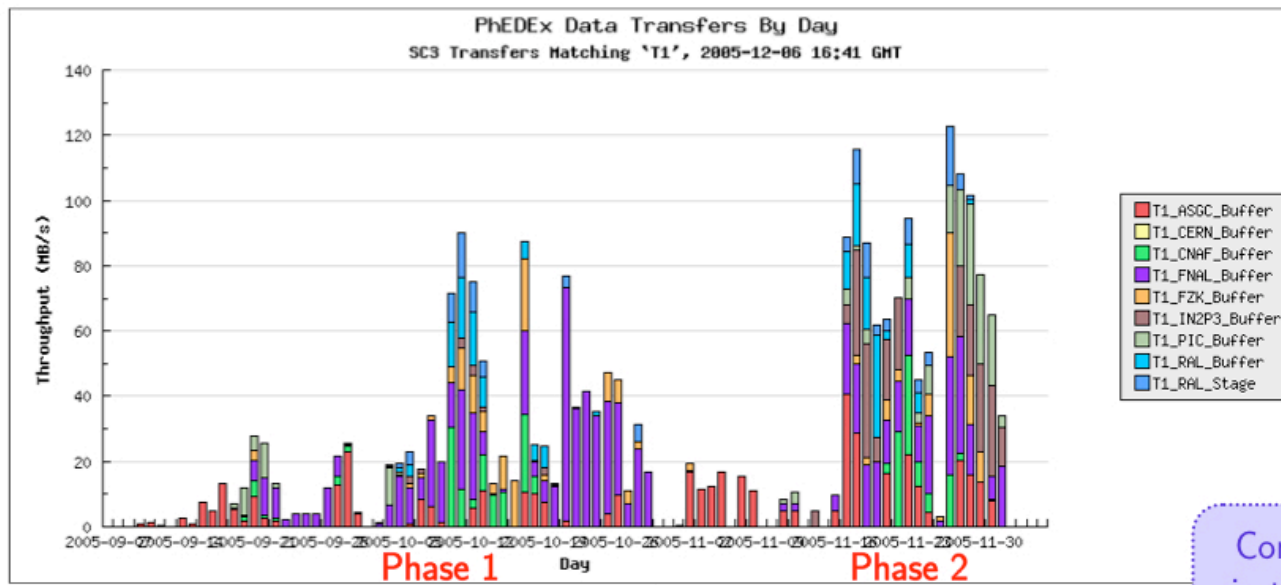
*) RAL result uncertain

Dec 7, 2005


CMS Week

5

Tier-1 WAN Transfers



Considerable improvement in transfer rate and quality for service phase 2.




Daily aggregate rate from Tier 0 to Tier 1s peaked at 120 MB/s.

Dec 7, 2005


CMS Week

7

Tier-2 WAN Transfers



Daily aggregate rate from Tier 1s to Tier 2s peaked at 45 MB/s, typically well below 30 MB/s.



Considerable improvement in quality, rate and system sustainability for service phase 2.

Dec 7, 2005

CMS Week

8

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 inefficiencies became apparent during challenge period
 - De-scoped to debugging pieces that did not work as expected.
- Lessons learned and principal concerns
 - Castor-2: Innumerable problems - 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



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

T1-T1 Channel Status

Key

T1 Site

- FTS Server Status
- Configuration of Channel Management

SARA

- FTS Server
- Manage Incoming Channels

RAL

- FTS Server
- Manage Incoming Channels

FZK


- FTS Server
- Manage Incoming Channels

IN2P3

- FTS server
- Manage Incoming Channels

CNAF

- FTS Server
- Manage Incoming Channels

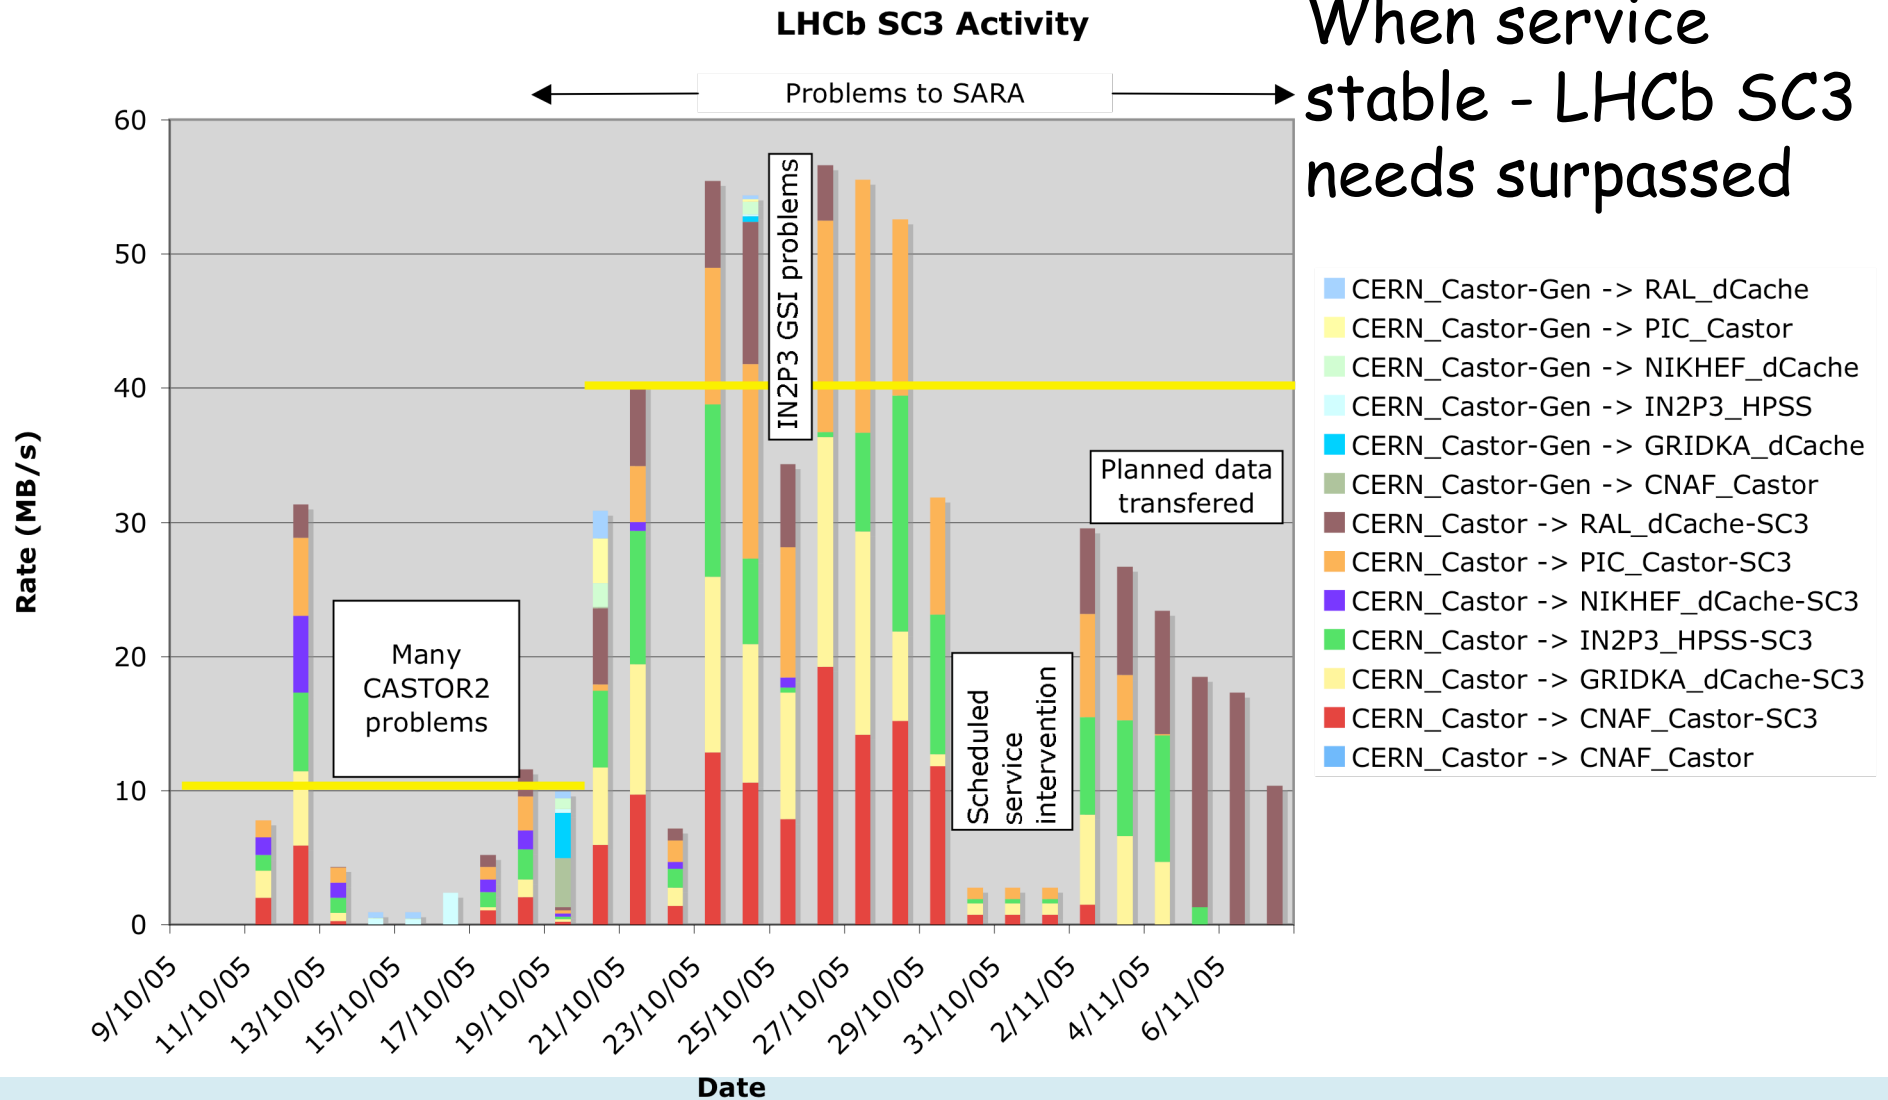


PIC

- No FTS Server
- Channels Managed by Source SE

FTS central service for managing T1-T1 matrix

Overview of SC3 activity



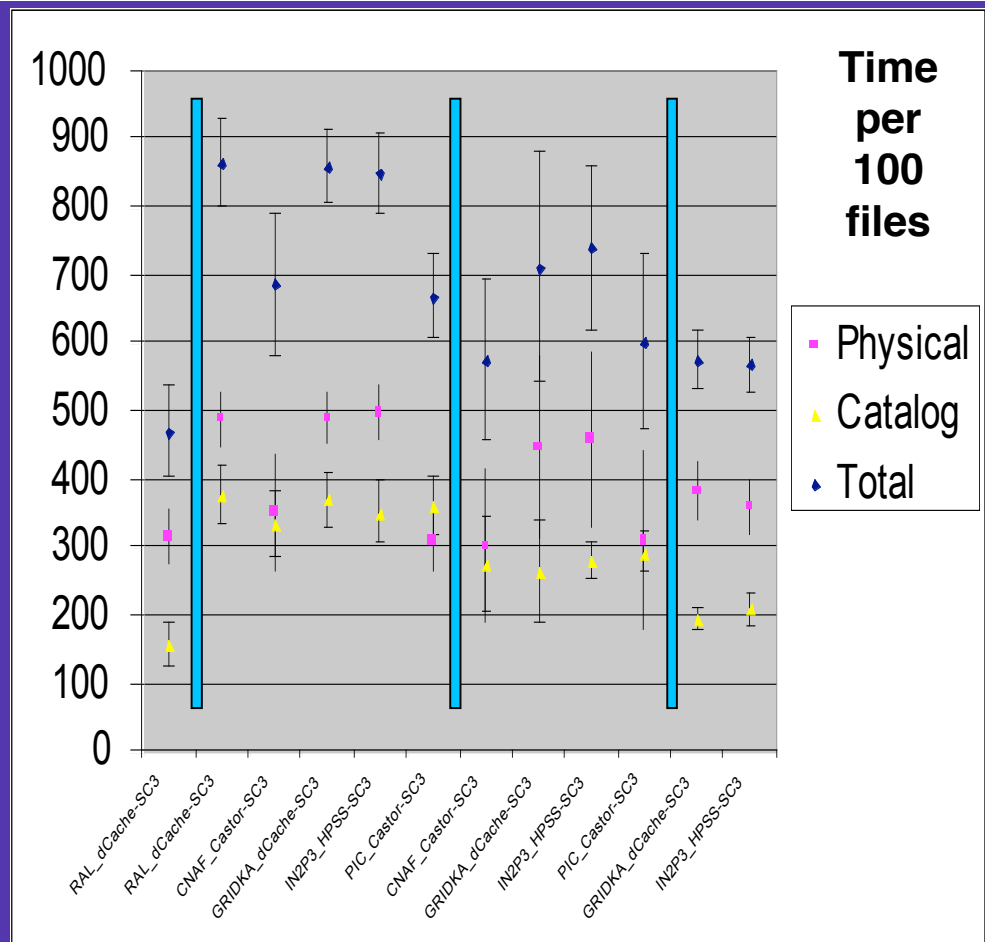
When service stable - LHCb SC3 needs surpassed

File removal

Four phases:

- RAL only
 - -RAL, CNAF, GRIDKA, IN2P3, PIC
 - -CNAF, GRIDKA, IN2P3, PIC
 - -GRIDKA, IN2P3
-
- 50k replicas removed in ~28 hours
 - 10k replicas at each site
 - Each site with its own agent
 - problems in the current middleware
 - Removing remote physical files is slow
 - Bulk removal operations are necessary
 - Different storage flavours are showing slightly different functionality :
 - No generic way to remove data on all the sites
 - LFC File catalog (un)registration operations are slow
 - Bulk operations within secure sessions are necessary

Secs →



Phase 1

1. Distribute stripped data Tier0 → Tier1's (1-week). 1TB
 - Succeeded but not a rate given in metrics of success
2. Distribute data Tier0 → Tier1's (2-week). 8TB
 - Succeeded - (nearly) achieved "acceptable" metric but not "success" metric
 - Like to repeat as part of the SC3 re-run
3. Removal of replicas (via LFN) from all Tier-1's
 - Failure to meet 24 hr metric Inconsistent behaviour of SRM, bulk operations needed, ...
4. Tier1 centre(s) to Tier0 and to other participating Tier1 centers
 - failure
 - FTS did not support third party transfer
 - Complicated T1-T1 matrix been set up - beginning to test

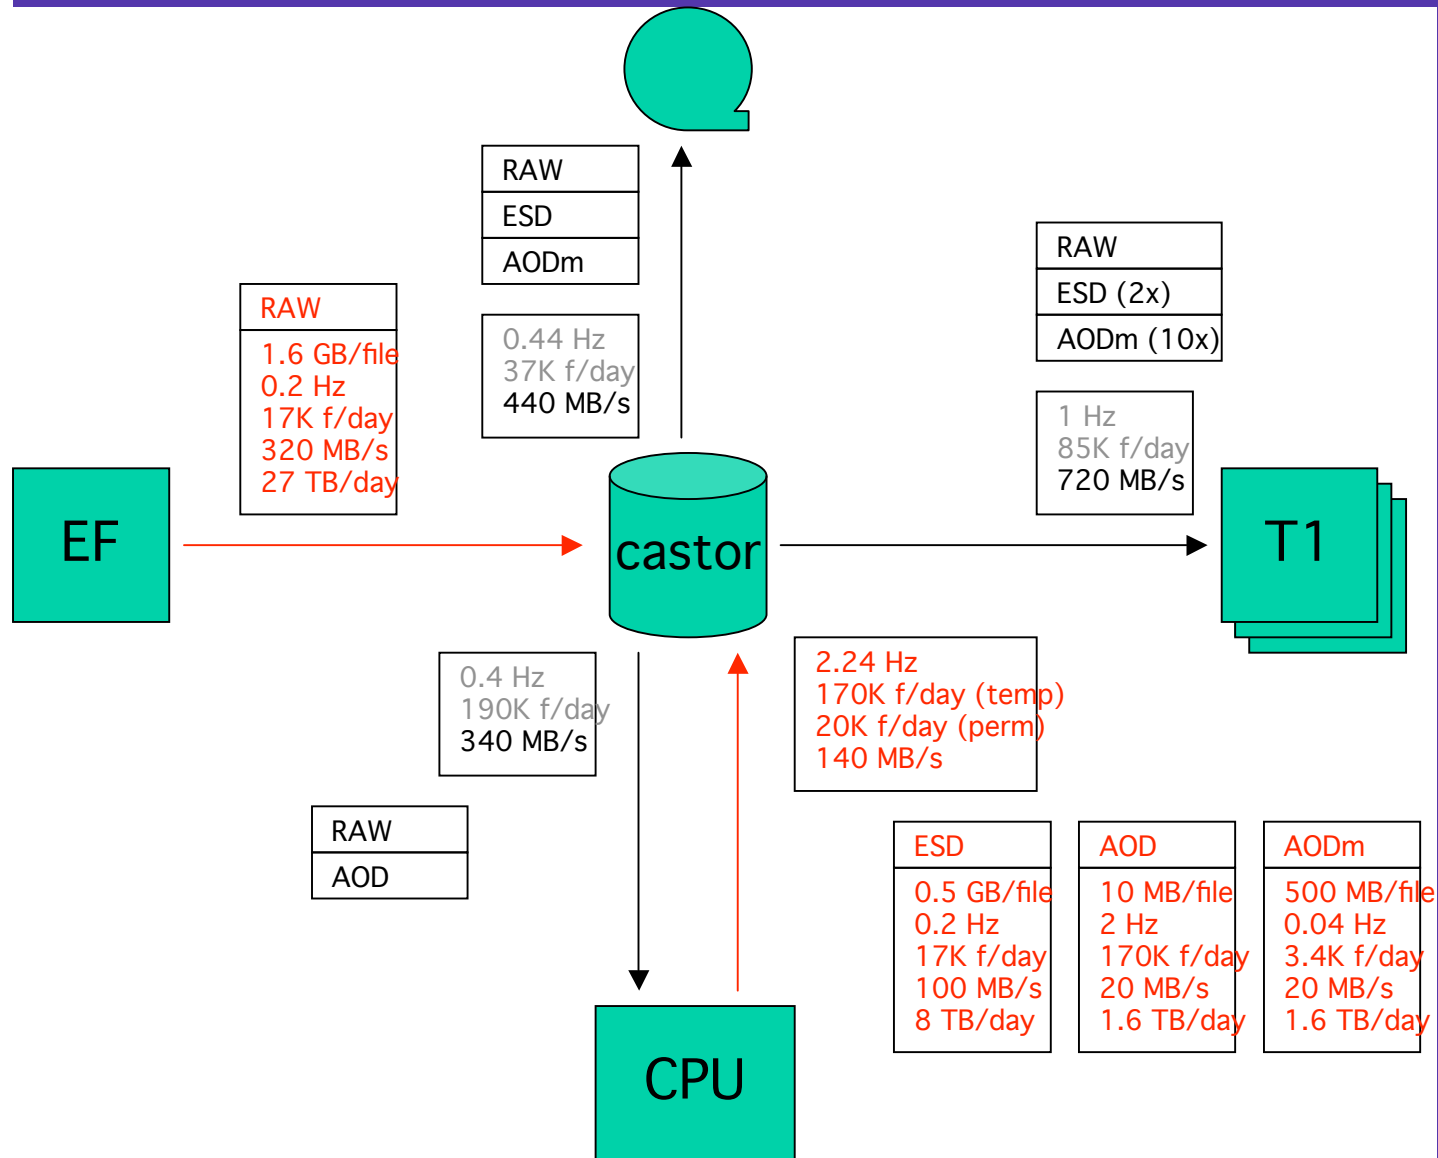
The ATLAS Experiment

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: 1-2 December



Transfers CERN - Tier 1 centres in the last 24 hours
Average throughput per hour

The current time at CERN is 18:15:00 Fri 02 Dec 2005




achieved quite good rate (sustaining >80 MB/s to sites)

Daily rates

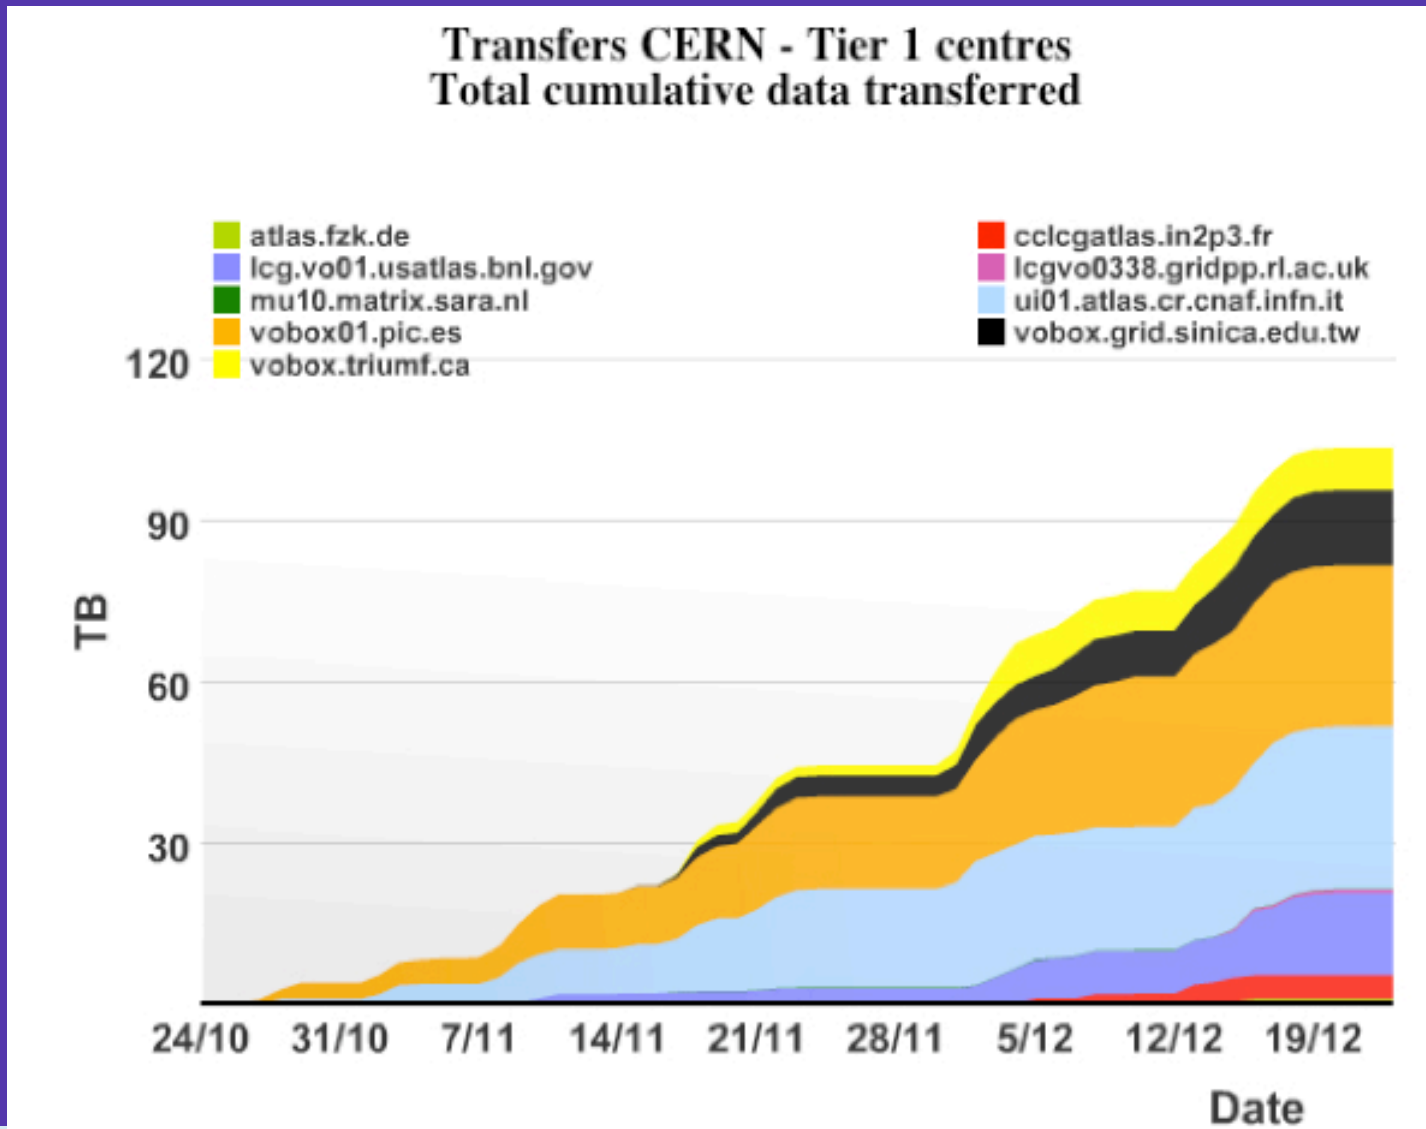


Transfers CERN - Tier 1 centres Average throughput per day



SC3 experience in 'production' phase

9/10 of ATLAS Tier1s used ...



General view of SC3



- ATLAS software seems to work as required
 - Most problems with integration of "Grid" and "storage" middleware (srm-dCache; srm-Castor) at the sites.
- Met throughput targets at various points
 - But not consistently sustained
- Need to improve communication with 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