

## SC3 - Experiments' Experiences

Nick Brook

In chronological order: ALICE CMS LHCb ATLAS



LCG France - Dec'05

# 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 (TO), 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

## <u>Status of production</u>

- 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 VOboxes 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:



- Total CPU work: 80 MSi2K hours; Total storage: 20 TB





## <u>Physics Data Challenge</u>



- 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

## <u>ALICE plans:</u>



- File replication with FTS:
  - FTS endpoints tested at all ALICE SC3 sites
  - Soon to start data migration, initially TO->T1
  - Test, if possible, migration Tx->Ty
- Re-processing of data with calibration at TO/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 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

## <u>Service Phase I Results</u>



Data

- 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 success rates in tables below

TO CHA

#### - 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

	VOIUITIE	Quality	110015	nate
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

Valuma

**Quality** = Successful transfers vs. those started **Hours** = Number of hours with successful transfers

= Volume / Hours

Rate



#### LCG France - Dec'05



LCG France - Dec'05

## <u>Summary of Experiences</u>



- 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





### 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

<u>LHCb Architecture</u> <u>for using FTS</u>

### 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



#### Phase 1



- Distribute stripped data Tier0  $\rightarrow$  Tier1's (1-week). 1TB
  - The goal is to demonstrate the basic tools
    - Precursor activity to eventual distributed analysis
- Distribute data Tier0  $\rightarrow$  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



#### Overview of SC3 activity



### File Deletion

- Initially (still) problematic as no single tool for removal
  - Implemented bulk operation of srm-adv-del
    - Physically removes file on dCache system (v1.6.6)
      - Older dCache version give different behaviour
    - Not complete removal from Castor systems
  - 50k replicas removed in ~28 hours
    - 10k replicas at five Tier1 sites
  - Physical removal time still high
    - ~4s per file was aim for 2s
  - Bottleneck at LFC also increasing operation time
    - SSL handshakes/authentication CPU intensive
    - max of 5 concurrent agents accessing

#### 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...



### 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
- 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





## <u>ATLAS SC3 goals</u>

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

# Concentrate on TierO 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 TierO



- 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)
  - TierO Management System (TOM)
  - Raw Data generator (Jerry)
  - Distributed Data Management (DDM) software (DQ2)

## Dataflow 2007



### <u>Snapshot of Activity</u> 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)





### SC3 experience in 'production' phase

- There are 10 Tier 1 sites have shipped data to 7 of them
  - It takes a lot of time and effort to set up the components necessary for each site
    - Usually human errors and inefficiencies
- The sites' components are not in general reliable
  - Storages are inaccessible, disks fill up, grid certificates expire, hardware/software upgrades etc etc...
- Castor @ CERN gave some problems...
- But the Grid middleware from LCG is stable and gives good enough performance
- In summary the data transfer needs constant babysitting - long way from fully functional Grids able to handle data rates.
- For details see
  - https://uimon.cern.ch/twiki/bin/view/Atlas/DDMServiceChallenge3Progress



## <u>General view of SC3</u>



- ATLAS have 2 weeks left of the exercise
  - The plan was to keep ramping up to 100% but this is not possible outside CERN with the current Grid TO-T1 infrastructure
    - try for 30% (210 MB/s throughput) or 40% (280 MB/s throughput)
- 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
  - The last 2 weeks or so we have been running on and off due to other constraints
- Need to improve communication with sites

#### <u>General Summary of SC3 experiences</u>

Extremely useful for shaking down sites, experiment systems & WLCG

- Many new components used for the 1<sup>st</sup> 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