



ALICE Quarterly Report 2008Q1

MONTE CARLO

Production

- Continuous MC production: 15 production cycles completed, including first iteration of MC for first physics
- Second iteration of MC for first physics pending (different detector misalignment)

Production	Description
PDC 08/LHC08x	p+p, charm, forced had.decays
PDC 08/LHC08w	p+p, beauty, forced had.charm decay, PYTHIA
PDC 08/LHC08v	jet-jet pp, PYTHIA, 15 GeV/c < Pt hard < 50 GeV/c
PDC 08/LHC08u	gamma-jet pp (2), PYTHIA, no quenching
PDC 08/LHC08t	MUON Cocktail pp, MB
PDC 08/LHC08s	p+p, beauty, with B->J/psi->ee decay
PDC 08/LHC08r	jet-jet pp, PYTHIA, Pt hard > 50 GeV/c
PDC 08/LHC08q	jet-jet pp, PYTHIA, hard > 100 GeV/c
PDC 08/LHC08p	gamma-jet pp, PYTHIA, quenching
PDC 08/LHC08b6	First physics pp, Phojet, No field, 10 TeV
PDC 08/LHC08b5	First physics pp, Pythia6, No field, 900GeV
PDC 08/LHC08b4	First physics pp, Pythia6, No field, 10TeV
PDC 08/LHC08b3	First physics pp, Phojet, 5kG, 10TeV
PDC 08/LHC08b2	First physics pp, Pythia6, 5kG, 900GeV
PDC 08/LHC08b1	First physics pp, Pythia6, 5kG, 10TeV

Storage

- 160 Mio MC events; 80 TB ESDs
- Output data carefully tuned to save storage
- Most of the data is available on T2 SEs for end users analysis

SE Name	AliEn name
1. Bari - dCache	ALICE::Bari::dCache
2. Catania - DPM	ALICE::Catania::DPM
3. CCIN2P3 - dCache	ALICE::CCIN2P3::dCache
4. CCIN2P3 - dCache_sink	ALICE::CCIN2P3::dCache_sink
5. CCIN2P3 - dCache_tape	ALICE::CCIN2P3::dCache_tape
6. CERN - C2PPS	ALICE::CERN::C2PPS
7. CERN - Castor2	ALICE::CERN::Castor2
8. CERN - Castor2X	ALICE::CERN::Castor2X
9. CERN - DPM	ALICE::CERN::DPM
10. CERN - se	ALICE::CERN::se
11. CNAF - Castor2	ALICE::CNAF::Castor2
12. CNAF - CASTOR2_sink	ALICE::CNAF::CASTOR2_sink
13. FZK - dCache	ALICE::FZK::dCache
14. FZK - dCache_sink	ALICE::FZK::dCache_sink
15. FZK - dCache_tape	ALICE::FZK::dCache_tape
16. GSI - dCache	ALICE::GSI::dCache
17. GSI - se	ALICE::GSI::se
18. GSI - se_tactical	ALICE::GSI::se_tactical
19. ISS - File	ALICE::ISS::File
20. JINR - dCache	ALICE::JINR::dCache
21. Legnaro - dCache	ALICE::Legnaro::dCache
22. NDGF - dcache	ALICE::NDGF::dcache
23. NDGF - dCache_sink	ALICE::NDGF::dCache_sink
24. NDGF - dCache_tape	ALICE::NDGF::dCache_tape
25. NIHAM - File	ALICE::NIHAM::File
26. Prague - Disk	ALICE::Prague::Disk
27. Prague - Disk2	ALICE::Prague::Disk2
28. RAL - Castor2	ALICE::RAL::Castor2
29. RAL - Castor2_sink	ALICE::RAL::Castor2_sink
30. SARA - dcache	ALICE::SARA::dcache
31. SARA - dCache_sink	ALICE::SARA::dCache_sink
32. SARA - dCache_tape	ALICE::SARA::dCache_tape
33. SPbSU - DPM	ALICE::SPbSU::DPM
34. Subatech - DPM	ALICE::Subatech::DPM
35. Torino - DPM	ALICE::Torino::DPM

Analysis

- End user analysis
 - Primary copy is at T2s SE and jobs run at the T2 holding the required data
 - Single user analysis jobs yield low CPU/Wall efficiency because I/O bound
- Analysis Train (organized analysis)
 - Many tasks processing a single data stream
 - Expect nominal CPU/Wall efficiency
- CAF analysis (PROOF and data local)
 - Data imported from Grid, stored locally on PROOF nodes
 - Performed on PWG choice data sets

December-February Cosmics run

RAW DATA

Data taking

- On line data replication to T1s **OK**
- On line first pass reconstruction **NOK**
 - Frequent changes in the code (first time AliRoot is confronted with real RAW data)
- On line condition parameters calculation (DAQ, HLT, DCS)
mostly OK
 - The conditions framework is fully operational
- On line reconstruction **OK**
- Monitoring and QA **partly OK**
 - Framework operational
 - Detector implementation in progress

Processing and Storage

- Data reconstructed offline after data taking
 - Pass 1 reconstruction processed at T0, data in CASTOR2@CERN (size ~20 TB)
 - 98% of data are reconstructible (total RAW size ~200 TB)
 - End user analysis performed by detector experts (Grid, local and CAF)
- Data replicated to T1s – 90% quasi online, during data taking, remainder offline
 - Pass 2 reconstruction being processed at T1s, data in MS@T1s, very careful evaluation of output data to store

ALIROOT

Status

- Strict release policy implemented
 - 2x version for cosmic reconstruction, every second week
 - Current version for MC production for first physics
 - [pp@0.9/10](#) TeV, w/wo B field
 - Detectors as installed
 - Store minimum needed for analysis
 - Size of ESD/AOD within Computing Model values
 - Open issues:
 - Raw data format not yet final
 - Aggressive optimization to keep memory footprint below 2GB/process (< 1.5 GB on 32 bit nodes and < 2.5 for 64 bit nodes)

Organized analysis

ANALYSIS TRAIN

Status

- Provides access to all analysis platforms with the same code
 - Usage: local, AliEn(Grid), CAF(PROOF)
 - Wagons (user code) provided by the PWGs
- Tested with large scale analysis of PDC08 data
- Optimization of PWG code ongoing with MC first physics data

CAF

Status

- Events are grouped in PROOF datasets which are generated from Grid file collections
- Files are staged from any Grid SE
- Users CPU and disk quota implemented
- As anticipated CAF/PROOF provides the look and feel of an interactive session
- Critical for 'fast track' analysis of RAW, calibration and ESDs
- Increasing load on the CAF from end users
- **We might have to close the CAF service if the PROOF support is not guaranteed anymore**

SERVICES

Status

- ALICE Grid services fully migrated to SLC4
 - 64 bit VO box in all sites with 64 bit WNs
- New AliEn version (v2-15) being deployed at all sites (67 in total)
 - Improvements in Package management, Job management (user quotas and priorities)
 - Secure access to storage fully implemented

Accounting

- ALICE resources accounting
 - Per process based – individual CPU types are taken into account
 - SI2K scaling factors taken from HEPiX benchmarking, discussed with site managers
 - Discrepancies noticed with the WLCG accounting in T1s and T2s
 - Discrepancies noticed with the WLCG accounting in T1
 - Pledges different from the ones at October C-RRB
 - Sites without VO box deliver resources !!
 - Delivered are systematically different
 - Harmonization is needed if not Scrutiny will turn nuts
 - ALICE trusts its accounting

Accounting

Accounting Tier 1. Sources: WLCG monthly report and ALICE MonALisa report, April 2008								
CPU								
Tier 1	WLCG T1 accounting			ALICE MonALisa			2008 C-RRB Pledges	
	Total Pledged	Delivered to ALICE (wall)	Fraction total	Pledged	Delivered Wall	Fraction	All	ALICE
	KSI2K	KSI2K	%	KSI2K	KSI2K	%	KSI2K	KSI2K
CERN Tier-0+CAF	8'083	1'668	46%	1602	801*	50%	15'851	2'300
CCIN2P3	1'733	506	52%	1060	387	37%	5'740	1'060
CNAF	1'475	49	41%	660	24	4%	3'000	660
FZK-GRIDKA	2'160	674	45%	600	365	61%	5'672	2'500
NDGF	906	149	41%	602	187	31%	2'172	1'102
NL LHC/Tier-1	2'014	41	13%	475	145	31%	4'382	317
RAL	1'505	25	41%	132	51	39%	3'139	132

* Does not include CAF

Accounting Tier 2. Sources: WLCG monthly report and ALICE MonALisa report, April 2008								
Tier 2	WLCG T2 accounting			ALICE MonALisa			2008 C-RRB Pledges	
	Total Pledged	Delivered to ALICE	Fraction total	Pledged	Delivered Wall	Fraction	All	ALICE
	KSI2K	KSI2K-hrs	%	KSI2K	KSI2K	%		
CZ-Prague-T2	164	28	55%	95	75	79%	170	95
FR-IN2P3-CC-T2	1500	122	165%	230			1500	230
FR-GRIF	1642	72	124%	205	79	39%	1642	205
FR-IN2P3-LPC	800	22	53%	240	0	0%	800	240
FR-IN2P3-SUBATECH	312	148	110%	434	182	42%	434	434
DE-GSI	660	34	12%	260	156	60%	660	660
HU-HGCC-T2	270	35	142%	90	58	64%	270	90
IN-DAE-KOLKATA-Tier2	450	5	3%	450	26	6%	450	450
INFN-BARI		18		68	12	18%		
INFN-CATANIA		104		307	206	67%		
INFN-LNL-2		27		68	35	51%		
INFN-TORINO		11		307	0.059	0%		
IT-ALICE-federation	750	160	49%	750	253	34%	750	750
KR-KISTI-T2	100	4					100	100
CYFRONET-IA64		5						
CYFRONET-LCG2		4		88	0	0%		
PSNC		314		88	224	255%		
PL-TIER2-WLCG	1250	314	129%				1250	265
RO-LCG	1300	272	74%	330	311	94%	1300	660
RU-Protvino-IHEP		50		64	34	53%		
IITEP		82		70	48	69%		
JINR-LCG2		3		155	7	5%		
RRC-KI		432		250	0	0%		
RU-SPbSU		5		3	22	733%		
ru-PNPI		68		77	125	162%		
RU-Troistsk-INR-LCG2		0		8	93	1163%		
RU-RDIG	3000	244		697	397	57%	3000	1050
UKI-LT2-Brunel		10						
UKI-LT2-IC-HEP		17						
UKI-LT2_IC-LeSC		6						
UKI-LT2-RHUL		1						
UK-London-Tier2	1017	34	33%	0	0		1017	0
UKI-NORTHGRID-LANCS-HEP		4						
UKI-NORTHGRID-LIV-HEP		6						
UKI-NORTHGRID-MAN-HEP		5						
UKI-NORTHGRID-SHEF-HEP		3						
UK-NorthGrid	313	18	266%	0	0		725	0
ScotGRID-Edinburgh		1						
UK-ScotGrid	313	1	49%	0	0		925	0
UKI-SOUTHGRID-BHAM-HEP		10		50	0	0%		
UKI-SOUTHGRID-BRIS-HEP		4						
UKI-SOUTHGRID-OX-HEP		5						
UKI-SOUTHGRID-RALPP		2						
UK-SouthGrid	959	27	107%				959	156
UA-	1166			1130	467	41%	1166	932

RESOURCES

Mass storage

- The most critical issue in 2008
 - Requirement 7.4 PB, Allocated 3.3 PB (CERN) + 2.98 PB (external)
 - Already used
 - at CERN 0.6 (MC) + 0.2 (Raw) + 0.2 (not Grid aware)
 - in external sites: 0.09 (MC) + 0.2 (Raw) + ?? (not Grid aware)
 - Potentially available: 2.3 PB (CERN) + 2.68 PB (ext)
 - Needed until end of the year for raw data and first physics MC: 3.6 PB
 - Margin of 1.3 PB, not sufficient for MC in preparation of run 2009
- We have a problem (anticipated and announced since years)
- Deleting data is not a solution

CPU and disk

- In April 2008,
 - only 30% of the pledged CPU power is available or usable
 - Only 22% of the requested disk capacity is available in WLCG SEs.
 - Out of which 22% is already used for MC data

Bandwidth T0 to T1s

- In 2007 we have requested a bandwidth of 60 MB/s which corresponds to a STYDT (C-TDR).
- In May 2008 the ALICE management has decided to take data in 2007 at the maximum rate allowed by the DAQ bandwidth: 300 MB/s
 - Can the T0 ALICE MS cope with this rate ?
 - Is there enough bandwidth left for export ?
 - Our pp data processing strategy might not work

Milestones

- MS-122 oct 07: FDR Phase II
 - Done
- MS-124 Feb. 08: Start of FDR Phase II
 - Done
- MS-125 Apr 08: Start of FDR Phase III & CCRC08
 - Postponed to June 2008 (delayed because of detector readiness)
- MS-126 Feb 08: Ready for CCRC 08
 - Done
- New milestones
 - MS-128 Jul 08: ready for data taking