ALICE-INT-2003-050

EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH
European Laboratory for Particle Physics

Internal Note
ALICE reference number
ALICE-INT-2002-
Institute reference number

Date of last change
26-08-2003

ALICE Data Challenge Planning
V 1.3
August 26, 2003

Authors:

P.Buncic, W.Carena, F.Carminati, R.Divià, A.Masoni, A.Morsch, F.Rademakers, K.Safarik, Y.Schutz, P.Vande Vyvre

Abstract:

This document contains the current planning for the ALICE Computing and Physics Challenges.

Content

Computing Data Challenges

Background

In 1997, the ALICE experiment and the IT division started the Mass Storage Project with the goal to assess the ALICE requirements in term of Mass Storage and evaluate the available products. Afterwards, the scope has been extended to cover a significant part of the functions of the ALICE DAQ and Offline projects at the CERN Tier 0 and the project became the ALICE Data Challenge. Active discussions are now under way to include also the HLT project in the tests.

In the framework of this project, since 1998, ALICE and the IT division have jointly executed several large-scale high-throughput distributed computing Data Challenges. The goals of these are to prototype the data acquisition and the offline computing systems, to verify their integration, to assess technologies and computing models, to test hardware and software components (including monitoring and HLT algorithms) of these systems in realistic conditions and to realise an early integration of the overall ALICE computing infrastructure.

Computing Data Challenge Planning

The ALICE Data Challenge milestones are listed in Table 1

Date	Mbytes/s	Tbytes to MSS	Nodes with Gbit Ethernet needed	Offline milestone
10/2002	200	200		Rootification of raw data -Raw data for TPC and ITS

9/2003	300	300	150	Integration of single detector HLT, at least for TPC and ITS - Quasi on-line reconstruction at CERN - Partial data replication to remote centres
5/2004	450	450	200	HLT prototype for all detectors that plan to use it - Remote reconstruction of partial data streams -Raw digits for barrel and MUON
5/2005	750	750	250	Prototype of the final HLT software Prototype of the final remote data replication (Raw digits for all detectors)
5/2006	750 (1250 if possible)	750 (1250 if possible)	300	Final test (Final system)

Table 1 ALICE Data Challenges milestones

The requirements in terms of bandwidth to the Mass Storage System are shown in Figure 1.

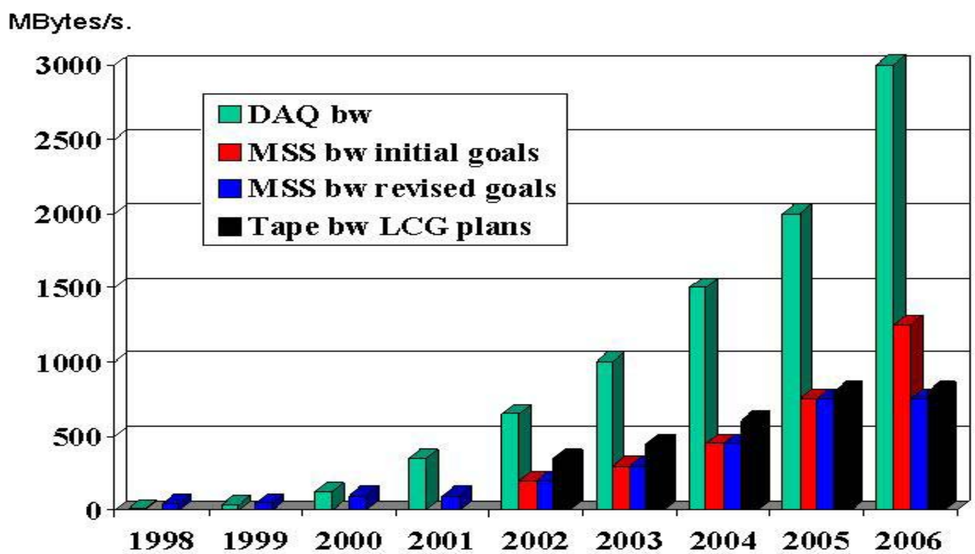

Figure 1 – ALICE DC bandwidth goals

The “DAQ bw” plot only concerns data acquisition software and network objectives. The "MSS bw initial goals" plot represents the ALICE initial goals for the bandwidth to mass storage system before the announced delay of LHC machine from 2006 to 2007, as they were presented to COCOTIME in June 2001. The "MSS bw revised goals" numbers represents the ALICE goals for the bandwidth to mass storage system, taking into account the LHC delay. The assumption is that the first heavy ion run will take place at the end of the first proton physics run, i.e. in March/April 2008 if there is no winter shutdown in 2007 like in the previous plan, or in November 2007 if there is a winter shutdown. The revised goals also take into account the document "LCG Prototype – Capacity and Guidelines for Scheduling Usage in 2002/3" by the PEB (27th of February 2002), documenting the agreement reached in the LCG Project Execution Board (PEB) concerning the capacity profile and scheduling guidelines for the LCG prototype equipment at CERN.

The model described in this document is plotted as "Tape bw LCG plans". It foresees to have 350MB/s in 2002 and 450MB/s in 2003. This gives to ALICE the margin to obtain the sustained target performances (200MB/s in 2002 and 300MB/s in 2003). The document describes the CERN prototype as sized to demonstrate at least 2/3 of the data recording rate required by ALICE during the first ion run. The capacity requirements are shown in Figure 2.

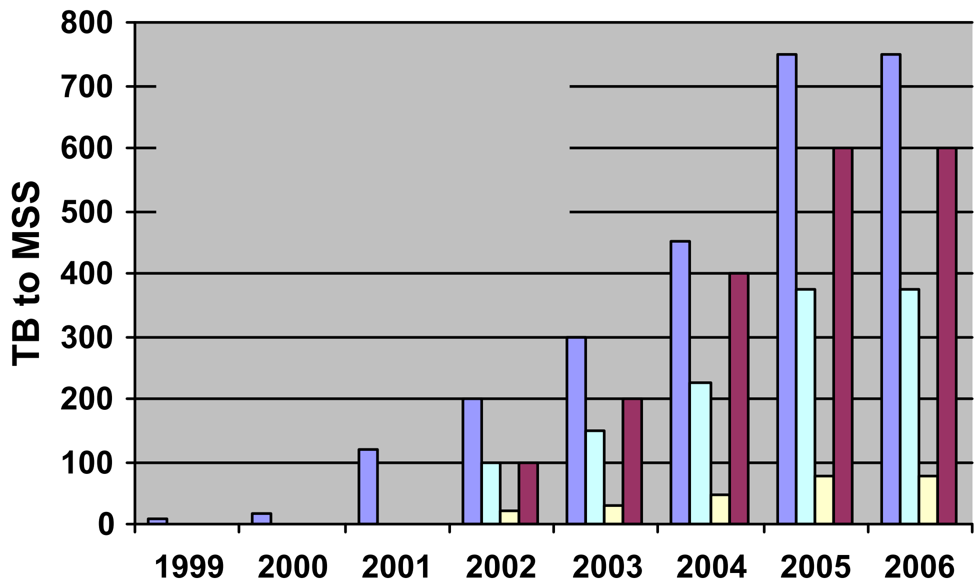

Figure 2 – Capacity requirements

Tapes will be recycled during the Computing Data Challenges, so we actually need about 50% of the total data to tape, as plotted in the "Tape storage need". Data will be used for testing purposes during three months after the Data Challenge, then they can be discarded; 10% of data however will be kept for a longer period, as indicated in the "Permanent storage need" plot. Table 2 summarizes the Computing Data Challenge needs each year. The period is usually agreed with IT Division.

	1998	1999	2000	2001	2002	2003	2004	2005
Total data to tape (TB)		7	15	120	200	300	450	750
Tape storage need (TB)					100	150	225	375
Permanent tape storage need (TB)					20	30	45	75
LCG plans capacity (TB)					100	200	400	600
DAQ bw (MB/s)	10	35	125	350	650	1000	1500	2000
MSS bw initial goals (MB/s)					200	300	450	750
MSS bw revised goals (MB/s)	45.9	46.9	100	100	200	300	450	750
Tape Bw LCG (MB/s)					350	450	600	800

Table 2 ALICE Computing Data Challenges Resources plan

ALICE Data Challenge III

The ALICE Data Challenge 3 (ADC III) was executed in 2001 and it included prototypes of the ALICE data acquisition system, the ALICE offline computing, the CERN Tier 0 system and the CASTOR Mass Storage System. The ALICE data acquisition software (DATE) was used for injecting simulated raw data of the ALICE TPC, produced with the AliRoot simulation program, in the data acquisition fabric, for event building and for controlling the data challenge. The data were then formatted with the ROOT I/O package. A data catalogue based on MySQL was established. The Mass Storage System CASTOR was used to store all the data. The components tested during the ADCIII are shown in Figure 3.

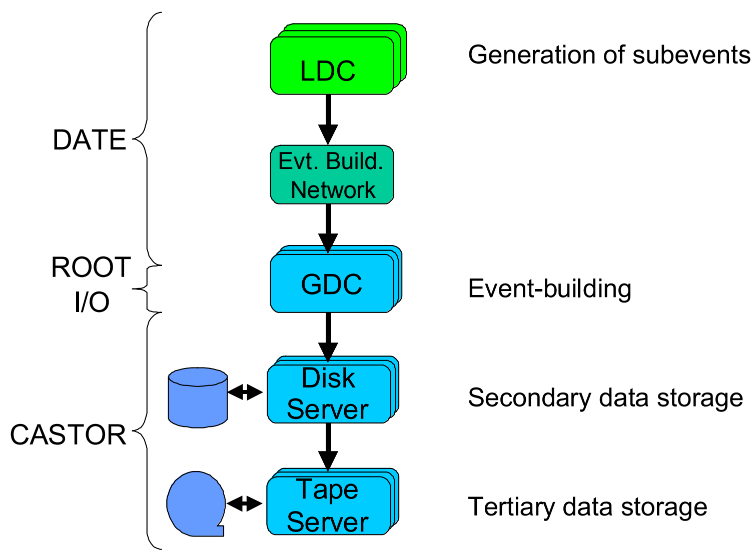

Figure 3 – Components tested during ADCIII

During ADC III, extensive measurements have been performed and major performance milestones have been achieved: up to 550 Mbytes/s in the computing fabric, up to 85 Mbytes/s through the whole chain from LDC to GDC during a week and more than 110 TBytes of data stored into mass storage. The fabric has proven to be stable and to deliver stable performance for periods of up to one week. The LHC computing testbed has been instrumental for the ADCs and ALICE plans to make further use of it for the larger tests scheduled in the future.

The Data Challenges are an opportunity to diagnose problems of implementation and scalability of all DAQ components, and of the Mass Storage System. During ADC III, the planned milestone of 100 MB/s sustained to tape has not been achieved. The two main causes were the limited performance of the generation of IDE-based disk servers used and the inefficient load balancing performed by CASTOR among the disk servers. Investigation of the performance problems is continuing. The ADC III key results are shown in Table 3.

Item	*Value*
Max throughput in DATE	550 Mbytes/s
Max throughput in DATE (ALICE-like traffic)	350 Mbytes/s
Max throughput in DATE+ROOT	240 Mbytes/s
Max throughput in DATE+ROOT+CASTOR	120 Mbytes/s
Average throughput during a week	85 Mbytes/s (> 50 TB/week)
Total amount of data through DATE	500 Tbytes in DAQ (2*10^7 events)
Total amount of data through DAQ+ROOT	200 Tbytes
Total amount of data through DATE+ROOT+CASTOR and recorded onto tape	110 Tbytes (100.000 files of 1 Gbytes)
Number of DATE runs	2200 runs
Longest DATE run	86 hours, 54 TB
Metadata database	10^5 entries

Table 3 – ADC III results

ALICE Data Challenge IV

The ALICE Data Challenge IV took place between June and December 2002. Its objective was a scalability test for the Data Acquisition system to control and handle hundred of nodes. The following performance targets were proposed:

- Data transfer inside the Data Acquisition system (minimum sustained throughput for a few hours):
    - 650 MB/s with ALICE detectors simulated data
    - 1 GB/s with flat traffic
- Data recording to Disk at 300 MB/s sustained.
- Data recording to Permanent Data Storage at 200 MB/s minimum sustained throughput for seven consecutive days.
- Around 200 TB of data being recorded to Permanent Data Storage.

The hosts used for the test were all SMP-based. The main production periods took place on the LCG testbed, based on boards equipped with dual Pentium III running at ~1 GHz. The hardware configuration of the ALICE Data Challenge IV is shown in Figure 4.

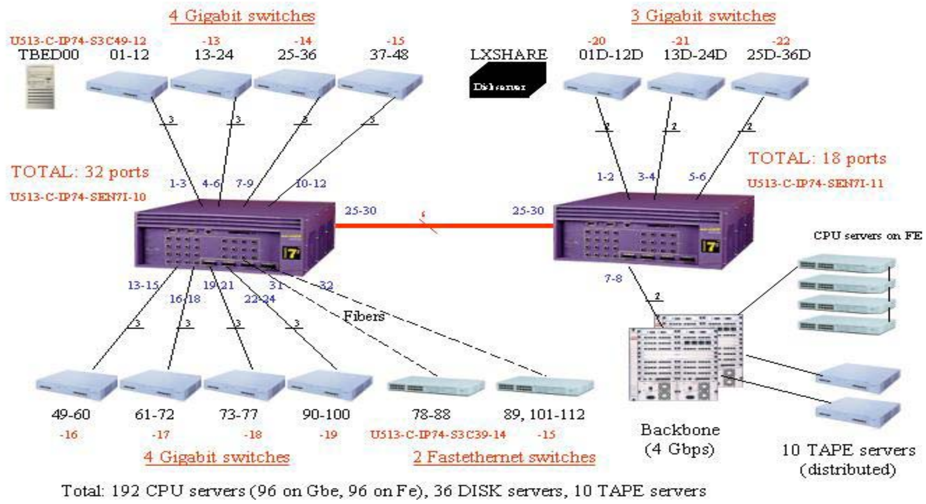

Figure 4 – ADC IV hardware configuration

ADC IV concluded in December 2002 with the following results:

- DAQ sustained bandwidth with flat data traffic of 1.8 GB/s (milestone was 1.0 GB/s), no recording, 5 days non-stop period.
- DAQ and data recording to disk using CASTOR with ALICE-like data traffic: data to disk sustained bandwidth 350 MB/S (milestone was 300 MB/s) and 140 TB of total data volume to disk.
- DAQ and data recording to EIDE-based disk servers (CASTOR), migration to STK 9940B (CASTOR) during 7 days non-stop: sustained bandwidth to tape of 280 MBytes/s (goal was 200) and 170 TB (goal was 200TB) of data volume to tape.

ALICE Data Challenge V

ALICE requirements for ADC V can be found in the 2003 COCOTIME ALICE requests under http://it-exp-coordination.web.cern.ch/it-exp-coordination/COCOTIME/2003/. The summary of the requests is reported in Table 4.

	Machines with Gbit/s Eth.	*Bandwidth to disk (MB/s)*	*Bandwidth to tape (MB/s)*
March 2003	*30*	*100*	
April 2003	*150*	*450*	*300*
October 2003	*30*	*100*	
November 2003	*150*	*450*	*300*

Table 4: Summary of the ALICE request for Data Challenge V

The following DAQ performance goals are proposed for the ADC V:

- System as in ALICE year 1 (2007, pp run and pilot PbPb run)

- 30 % of final system (100 LDC's);
- Scalability up to 150 nodes (LDC's + GDC's);
- Increase performances:
    - From DAQ to tape: 300 MB/s sustained over one week
    - From DAQ to disk: 450 MB/s peak
    - Inside DAQ: 3 GB/s

The following technology objectives are proposed:

- Network: New generation of NIC cards (Intel Pro 1000) and 10 Gbit Eth. Backbone.
- Storage: New IDE-based disk servers (RFIO @ 90 MB/s) and STK9940B tapes (~ 30 MB/s, ~ 200 GB/vol).

The ALICE GDC software runs both on 32 and 64 bit platforms. ALICE is interested in using 64 bit machines during the ADC V period in order to perform comparison on the price/performance ratios between the two types of platforms.

Monitoring during ADC V

During the ADC V, a new functionality will be introduced and its performance tested. The objective is to provide means to the remote sites to monitor data produced at CERN. Two mode of operation will be tested. In the first one the data stream is not modified and data are stored in the CASTOR system and registered in the AliEn data catalog. Remote sites request data without selection from the data catalog; copy data locally on a temporary storage, process them and send the result (in the simplest case in the form of a display) back to the production site. In the second mode each remote site is specialized in a particular processing requiring a well-defined sub-set of events (characterized for example by the trigger pattern).

A typical application of such a mode would be the calibration of a given detector. To that purpose, in addition to the regular data flow that remains unchanged, data will be sampled according to a pre-defined sampling fraction, and re-clustered into files containing only events of one type. The file is registered in a new data catalog with an appropriate tag and stored in the castor system using a server different from the one used for the main stream. Since only a fraction of events will be re-clustered, this will have only a minor impact on the storage needs. Remote sites can then request the events of interest from this new catalog, and get the physical file from the dedicated CASTOR server, improving, thus, the efficiency the system.

Physics Data Challenges

AliRoot and AliEn

The cornerstones of the ALICE distributed computing model are AliRoot [1], an Object Oriented framework written in C++ and AliEn [2], the production environment that implements several components of the Grid paradigm needed to simulate, reconstruct and analyse data on distributed computing resources. AliRoot uses directly the ROOT [3] framework that offers the needed performance and functionality in a simple and user-friendly way. ROOT provides data persistency at the file level, an interface to various utility libraries, visualization, a graphical user interface, virtual Monte Carlo and a geometrical modeller. This approach allows fast prototyping and development by using local files containing one or more events, which is very much appreciated by ALICE users and developers. We have extended the capabilities of ROOT by providing an AliEn specific implementation of the abstract TGrid class to allow ALICE users a transparent access to resources on the Grid as if they were local. From the ROOT prompt, users can authenticate themselves; access distributed datasets or request the execution of their algorithms across distributed datasets and retrieve the resulting objects in an interactive or batch session.

AliEn provides the two key elements important for large-scale distributed data processing – a global file system [4] for data storage and the possibility to execute jobs in a distributed environment, harvesting available CPU resources in an optimal way. The system takes care of job splitting and execution. When required, the datasets can be replicated and cached and AliEn and tries to optimize network traffic using scheduled file transfer and the hierarchy of file caches. The result of each job is optionally validated and the final output from many concurrent jobs is merged together and presented to the user as a dataset in his portion (directory) of the global logical file namespace (AliEn File Catalogue). The system keeps track of the basic provenance of each executed job or file transfer.

Objectives

The goals of the Physics Data Challenges (PDC) are:

- Determine readiness of the offline framework for data processing;
- Validate the distributed computing model.

In the current ALICE computing model, the CERN Tier-0/Tier-1 complex will be responsible for the primary data acquisition function and for the first-pass event reconstruction. Subsequent reconstructions will be shared among the different Tier-1 centres that will therefore collectively store a copy of the raw data. This model is very similar to the one of the other LHC experiments, and leads to the requirement of having quite a substantial amount of computing resources at CERN, and it is ALICE's baseline requirement. However we believe we must also consider the case in where it might be difficult to provide these resources concentrated at CERN. Therefore we have evaluated a backup model where the reconstruction charge is shared among all Tier-1s, including the one at CERN.

Table 7 shows the computing requirements both in the distributed reconstruction scenario and in the baseline scenario. The differences between the two models depend mainly on the computing capacity available at CERN (and, to a lesser extent, on the standard Tier-1 capacity). ALICE requirements at CERN, in the baseline scenario are the same as in Task Force 1 report and have been confirmed during the LCG Phase 2 recosting procedure. The requirements in the *distributed reconstruction* scenario represent only an alternative hypothesis and do not constitute a change in the current ALICE requirements at CERN.

One of the main goals of the ALICE PDC-3 will be to provide the necessary information to an appropriate evaluation of the two models. The PDC milestones are reported in the Table 5. Percentages refer to the amount of simulation that we will perform during data taking (= 1/10 of the simulated data in one year at standard conditions)

Period (milestone)	*Fraction of the final capacity (%)*	*Physics Objective*
06/01-12/01	1%	pp studies, reconstruction of TPC and ITS
06/02-12/02	5%	First test of the complete chain from simulation to reconstruction for the PPR Simple analysis tools. Digits in ROOT format.
01/04-06/04	10%	Complete chain used for trigger studies. Prototype of the analysis tools. Comparison with parameterised MonteCarlo. Simulated raw data.
01/06-06/06	20%	Test of the final system for reconstruction and analysis.

Table 5 Physics Data Challenge milestones

The estimates are assuming the target of simulating 5%, 10% and 20% of the events simulated in 1 year of standard data taking.

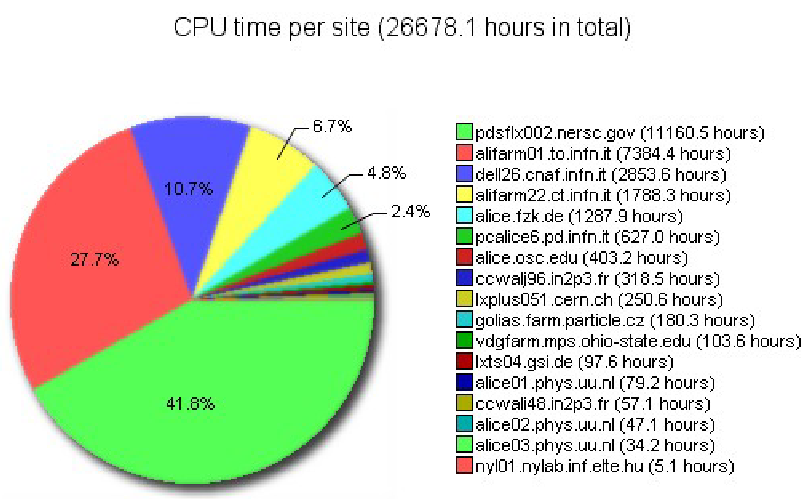

Pb-Pb production

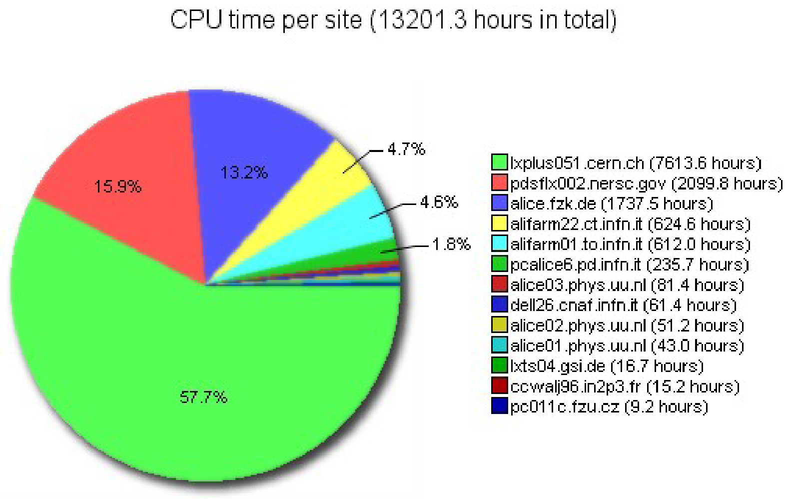

p-p production minimum bias

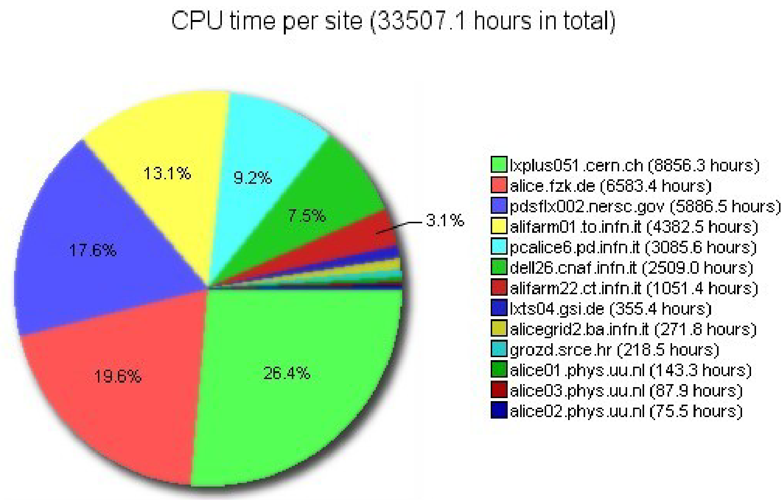

p-p production hadronic charm

Figure 5 Geographical distributions of ALICE Physics Data Challenge 2 production CPU time (2002)

Figure 5 shows the geographical distribution of the CPU usage during the 2002 PDC. It provides the fundamental inputs for the evaluation of the next PDC needs and constitutes a basic experience for the test of large scale distributed production. All the productions have been made using AliEn and were distributed across several centres in EU, US and elsewhere.

Physics Data Challenge III

Physics requirements

Recently the interests in heavy-ion physics has been significantly shifted towards hard parton scattering both due to development in theory and the data coming from RHIC experiments, which may show the evidence for jet quenching. During recent years, new sub-detectors have been added to the ALICE experiment design, substantially improving its capability to detect high-momentum particles. In order to study the ALICE performance in jet physics, we have to modify our strategy for event simulations and mixing of the signal with the underlying events. Practically in each central Pb-Pb collision (underlying event) a 20 GeV jet is produced in addition to hundreds of mini-jets above 5 GeV inside the ALICE acceptance. This means to study the quenching of jets up to few tens of GeV, where the relative energy loss is most significant, and where we cannot use the mixing of a signal with an underlying event. For harder jets (of the order of 100 GeV) we will still use the mixing, however, due to the fact that the transverse energy flow in underlying events will have significant local fluctuations, we have reduced by one order of magnitude the number of times one event could be reused (from 500 times, which was considered for other probes, to about 50). Preliminary estimates show that to study the jet quenching in the region around 20 GeV using particle correlations, 10^5 Pb-Pb collisions simulated with the HIJING heavy-ion event generator will be needed. The same events can be used with mixed-in signals of 100 – 200 GeV jets in order to investigate the performance of ALICE detector in all accessible p_t range. For all other observables this amount of simulated events would be more than sufficient.

Resource Estimate

Two main elements entered in the evaluation of the storage and CPU needs:

1. The current knowledge of the computing needs derived from our Data Challenges experience;
2. The estimated needs, for simulation, at the time of detector full operation in p-p and Pb-Pb runs in 2008.

On this basis, PDCs are planned scaling as 1%, 5%, 10% and 20% of the estimated final need for simulation.

The capacity for DC3, scheduled to start at the beginning of 2004, has been evaluated in the way explained below. The data for the event size and CPU power were derived from the previous PDC and are reported in Table 6.

Cpu to simulate 1 event	kSI2k*s	1,E+03	108,0	18000
Cpu to reconstruct 1 event	kSI2k*s	1,E+03	3,6	900
Size RAW simulated event	MB	1,E+06	0,8	600
Size ESD simulated event	MB	1,E+06	0,02	7

Table 6 ALICE parameters for Physics Data Challenge 3

As it can be seen from Table 6, both the size and the CPU power needed for a simulated Pb-Pb event, imply the utilisation of large amount of resources. To reduce somehow the usage of resources, the following technique has been developed and successfully applied to the previous PDCs. A reasonably reduced set of central Pb-Pb events is generated. Then each of these events (underlying events) can be combined with signal events and reused with an appropriate mixing procedure. In the last PDC each of the underlying events was used up to 500 times. However the focus of the next PDC will be shifted toward the study of hard processes such as jet studies. One of the main problems is the fluctuations in the background events that might hide low energy (100-

200GeV) jets. In this case, not to introduce a systematic bias, the "background event" can only be reused 50 times at most, depending on the energy domain of the jets.

Signal generation, mixing and reconstruction are extremely fast procedures and introduce a negligible contribution to the total amount of required resources. Therefore they are performed on the fly. On the basis of these estimates, the following resources plan is derived for PDC 3 (10 % scale of the simulation at standard conditions). On the basis of the current experience a development margin or inefficiency of 50% for the storage and 20% for CPU is accounted.

The production will be performed using AliEn. LCG will be one Computing Element of AliEn. This way LCG resources can be seamlessly integrated. ALICE aims at using LCG resources as much as possible and strongly relies on them.

The PDC spans over two quarters and is scheduled during the first quarter of 2004. The simulation production starts in the first quarter. At the same time reconstruction tests and analysis are performed on reduced data sets. The simulated data are produced at all available centres and then transferred (at least partially) to CERN. The CERN share in the simulation production phase will depend on the resources availability. In this model CERN contribution is evaluated at the level of an average Tier-1, i.e. about 15%.

Bandwidth requirements

About 90 TB of data will be shipped to CERN over a period of about 2 months (month 2 and month 3 of 2004). This corresponds to an average bandwidth to CERN of 160 Mb/s with a peak of 190 Mb/s during month 2. Assuming a dedicated CPU capacity for the data challenge of 186 KSI2K at each Tier-1 the data volume generated by simulation at each Tier-1 during the peak period (month 2, 2004) corresponds to 80 Mb/s.

The corresponding data traffic in the reconstruction phase (second quarter 2004) is estimated at about 80 Mb/s (sustained average bandwidth over three months from CERN to outside centres)

The bandwidth needs during analysis are not taken into account at this stage. The analysis system, based on Root/PROOF/AliEn, has been devised to minimise the data transfers. The effective performances depends however on the availability of resources where the data are located. We expect to obtain valuable information on this aspect from the data challenge results.

Having the data at CERN is justified by the fact that the second phase will simulate a situation comparable to real data taking period when all the data originates from CERN. The goal of the Data Challenge is to test not only distributed analysis but distributed reconstruction as well.

In this context too, the CERN share, from the CPU charge point of view, has been estimated at the level of an average Tier-1. Reconstruction (and analysis) will enter the production phase during the second quarter 2004. The Data Challenge is limited to a six months period, however the data will be used for physics analysis also during the rest of the year. This will have a limited impact on the CPU needs (that will stay within the standard baseline) but will imply storage needs for a longer period.

Figure 6 provides the evolution of CPU and storage requirements during the six months of the Data Challenge period according to the schema presented above.

Thanks to the mixing technique, simulation needs are relatively reduced with respect to reconstruction needs where the full statistics is involved. In this model, a portion (possibly all) of the simulated RAW data are produced remotely, copied to CERN and then deleted from the original location. The data are then replicated back to the participating Tier1's for reconstruction. The ESD are partially replicated at each participating Tier-1. A small fraction is replicated at Tier-2s also.

The model makes no distinction between the Tier0/Tier1 complex at CERN and the other Tier-1 centres for reconstruction. This allows true distributed production according to the availability of resources. This approach allows ALICE to benefit from all Grid technology developments without forcing external dependences. CERN plays its Tier-0 role as origin of the RAW data.

A relevant point concerns the Tier-2 role. At present the gained experience enabled to improve the evaluation of the Tier-2 role, which was not accounted previously. This is the first effort of calculating their impact in the model.

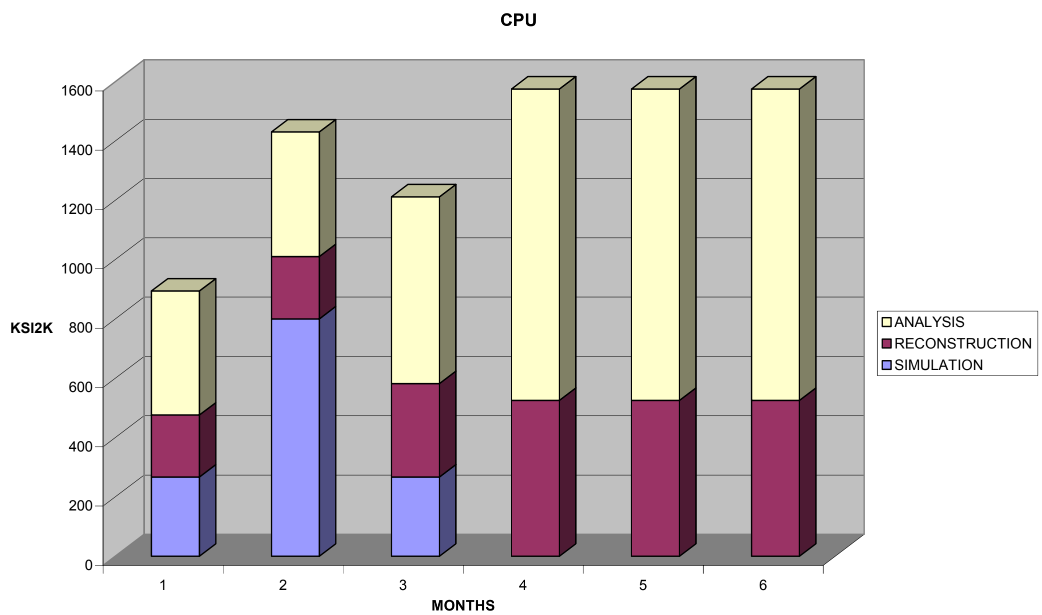

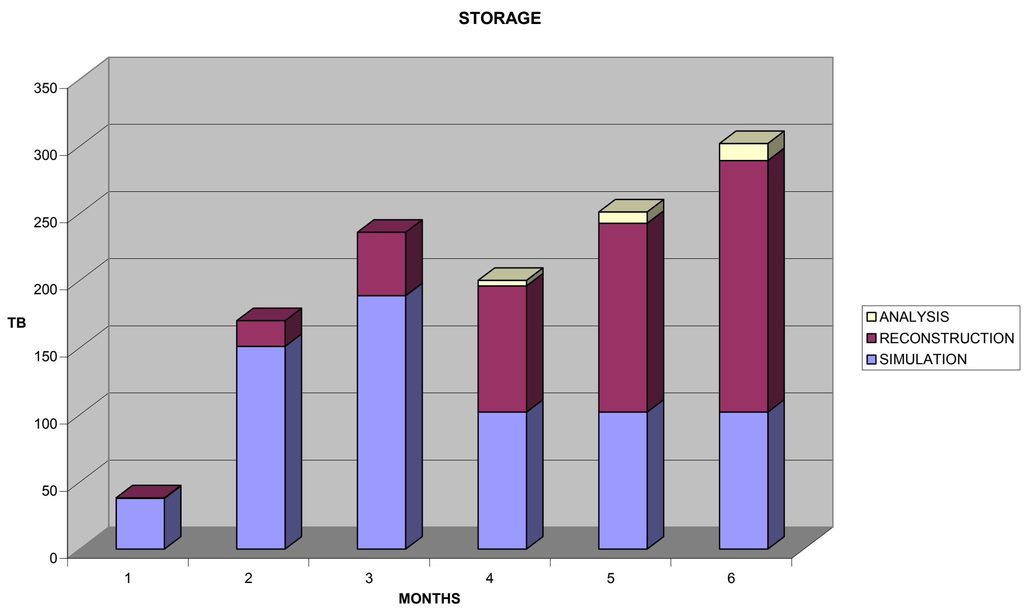

Figure 6: evolution of the CPU and storage requirements during the six months of the data challenge period

Table 6 contains CPU and storage requirements for PDC III.

These estimates are based on the resources strictly needed for the Physics Challenges. Code development special studies and Grid experimentation are not taken into account. System inefficiencies, such as downtimes or other operational problems, are not included either.

A load balance hypothesis where CERN contributes for 15% of the total CPU capacity is also shown. Actually, as explained above, the ALICE model has no specific constraint on the CPU load balance among the regional centres. Here CERN gives a contribution of 15% (the same level as the other Tier-1 centres).

The storage capacity at the different centres takes into account the different centres tasks. In particular since PDC aims at reproducing, on reduced scale, the real data taking environment, CERN must be the origin point for the real data.

The ALICE code (simulation, reconstruction, analysis) runs both on 32 and 64 bit platforms. The 64 bit machines are more expensive but rate higher I/O performances and can represent one of the future solution for HEP computing. It is therefore interest of ALICE to use clusters of 64 bit machines during the Physics DC-3 period in order to perform a test and comparison, in a realistic case, of the two environments.

Year,Quarter	04Q1	04Q2
CPU	KSI2K x 3 months	KSI2K x 3 months
Total Requirement for Simulation	444	
Total Requirement for reconstruction	303	467
Total Requirement for analysis	606	933
Total CPU Requirement	1354	1400
STORAGE	TB	TB
Total Requirement for Simulation	188,7	102
Total Requirement for reconstruction	63	187
Total Requirement for analysis		13
Total Storage requirements	251	301
LOAD BALANCE CPU	KSI2K x 3 months	KSI2K x 3 months
CERN	203	210
TIER-1s OUT CERN	744	770
Tier-2s	406	420
Total CPU Requirement (check)	1354	1400
LOAD BALANCE STORAGE	TB	TB
CERN	109	143
TIER-1s OUT CERN	93	111
Tier-2s	49	48
Total Storage requirements (check)	251	301

TYPICAL TIER-1 REQUIREMENTS		
CPU	KSI2K x 3 months	KSI2K x 3 months
	186	193
STORAGE	TB	TB
	23	28

Table 6: Requirements for CPU and Storage in 2004 for Physics Data Challenge

Physics Data Challenge IV

Detailed planning for the PDC IV has not yet been done. The *resource budget* available is twice that of Data Challenge III. In Table 7 we have indicated these resources, within which we will have to accommodate our physics needs. These will be defined during 2005 from the experience of PDC III based on the current views on the ALICE Physics programme.

Year,Quarter	06Q1	06Q2
CPU	KSI2K x 3 months	KSI2K x 3 months
Total Requirement for Simulation	889	
Total Requirement for reconstruction	536	933
Total Requirement for analysis	1073	1867
Total CPU Requirement	2498	2800
STORAGE	TB	TB
Total Requirement for Simulation	377	204
Total Requirement for reconstruction	125	374
Total Requirement for analysis		25
Total Storage requirements	502	603

Table 7: Estimated resources for PDC IV

Computing Resources planned by ALICE in standard data taking year

Parameter	Unit	ALICE REQUIREMENTS (DISTRIBUTED RECONSTRUCTION SCENARIO) ALICE		ALICE REQUIREMENTS (BASELINE SCENARIO) ALICE	
		p-p	Pb-Pb	p-p	Pb-Pb
Nb of assumed Tier1 not at CERN		4			
Event recording rate	Hz	100	50	100	50
RAW Event size	MB	0,2	25	0,2	25
REC/ESD Event size	MB	0,02	2,5	0,02	2,5
AOD Event size	MB	0,005	0,25	0,005	0,25
TAG Event size	MB	0,001	0,01	0,001	0,01
Running time per year	M seconds	10	1	10	1
Computing time per year	G seconds	0,02		0,02	
Events/year	Giga	1,00	0,05	1	0,05
Storage for raw data	PB	0,2	1,3	0,2	1,3
Storage for real (rec+AOD+TAG) data	PB	0,08	0,4		
Total Storage for real (raw+rec+AOD+TAG) data	PB	0,28	1,6	0,28	1,6
RAW SIM Event size	MB	0,4	600	0,4	600
REC/ESD SIM Event size	MB	0,04	5	0,04	5
Events SIM/year	Giga	0,1	0,0001	0,1	0,0001
Number of rec passes	Nb	2	2	2	2
Storage for simul. data	PB	0,05	0,69	0,05	0,05
Tape storage at CERN TO	PB	1,5		1,9	
Tape storage at CERN T1	PB	0,3		0,2	
Tape storage at each Tier1 (Avg)	PB	0,5		0,2	
Total Tape storage / year	**PB**	**3,9**		**2,9**	
Total bw at CERN T0	MB/s	248	1250	248	1250
Tape bandwidth at each Tier1 (Avg)	MB/s	16	182	3	14
Disk storage at CERN	PB	0,19		0,4	
Disk storage at each Tier1 (Avg)	PB	0,22		0,2	
Total Disk storage	**PB**	**1,29**		**1,1**	
Time to reconstruct 1 event	kSI-2k sec	3,6	900	3,6	900
Time to simulate 1 event	kSI-2k sec	27	20250	27	20250
CPU for 1 rec. pass/y (real data) in six months	kSI-2k	361	4509	361	4509
CPU for 1 SIM pass/y (simul+rec) in six months	kSI-2k	153	2356	153	2356
TOTAL CPU REAL DATA reconstruction, calib	kSI-2k	586	4779	586	4779
CPU simulation	kSI-2k	153	2356	153	2356
CPU analysis	kSI-2k	7875		7875	
Total CPU at CERN	kSI-2k	2362		7402	
Total CPU each Tier1 (Avg)		3347		2087	
Total CPU	**kSI-2k**	**15750**		**15750**	
WAN, Bandwidths					
Tier0 - Tier1 link, 1 expt	Mb/s	2000		1500	
Tier1 - Tier2 link	Mb/s	622		622	

Table 7: Computing Resources for standard data taking year

[1] R. Brun, F. Rademakers, ROOT – An Object Oriented Data Analysis Framework, Nucl. Instr. and Meth. A 389 (1997) 81, http://root.cern.ch
[2] P. Buncic, A. Peters, P.Saiz , The AliEn system, status and perspectives, http://arXiv.org/pdf/cs.dc/0306067
[3] R. Brun, P. Buncic, F.Carminati, A.Morsch, F.Rademakers, K.Safarik.: Computing in ALICE, Nucl. Instr. and Meth. A 502 (2003) 339-346
[4] A. Peters, P. Buncic, P.Saiz , The AliEn system, status and perspective, http://arXiv.org/pdf/cs.dc/0306067