



LHC Computing Grid Project

Quarterly Status and Progress Reports

2005 Q4

7 March 2006

Table of Contents

WLCG High Level Milestones.....	3
---------------------------------	---

Grid Sites

1. ASGC	5
2. CC-IN2P3	7
3. CERN	9
4. FZK	11
5. INFN	13
6. NDGF	15
7. PIC	17
8. RAL	19
9. SARA-NIKHEF	21
10. TRIUMF	23
11. US ATLAS	25
12. US CMS.....	27

Areas and Projects

1. Applications Area.....	29
2. Deployment Area	33
3. ARDA.....	35
4. Distributed Databases	37
5. Grid Deployment Board.....	39

Experiments

1. ALICE	41
2. ATLAS	43
3. CMS.....	45
4. LHCb.....	47

QUARTERLY STATUS REPORT					
Project Name			Date		
WLCG			2-Feb-2006		
Report Period			Author Name		
2005 Q4			Les Robertson		
Milestones for the Quarter					
ID	Date	Milestone	Coord.	Status	Comments
SC3-1	01.09.05	Service Challenge 3: start of stable service phase	J.Shiers	Done	Including at least 9 Tier-1 and 10 Tier-2 sites.
SC3-2	31.12.05	Service Challenge 3: successful completion of stable service phase 5 Tier-1s and 5 Tier-2s must have achieved the following targets: (a) appropriate baseline services operational (b) availability better than 80% of the levels specified in Annex 3 of the WLCG MoU (adjusted for sites that do not provide a 24 hour service) (c) Success rate of standard application test jobs greater than 80%(excluding failures due to the applications environment and non-availability of sites)	J.Shiers	Done - but throughp ut tests to be re-run in Jan/Feb 2006	All Tier-1s and ~20 Tier-2s active (a) appropriate baseline services operational at all sites (note that, as planned, VO Boxes were in operation for test and evaluation at a subset of sites) (b) measurement tools were being developed during SC3; using the basic SFT tests in the EGEE grid as measure of availability, by November average availability was better than 80% for ALL sites in the EGEE grid (not just the LCG SC3 sites) on a 24 hour basis (c) There are not yet standard test jobs for all experiments, and there is no consistent measure of job failure rate. However, taking all jobs submitted, ALICE, ATLAS and LHCb report grid/system failure rates of less than 20%, and analysis of the EGEE Resource Broker logs (by D.Colling /Imperial) shows for these experiments significantly less than 20% grid-visible failure rates.
OPN-1	31.12.05	Tier-0/1 high-performance network operational at CERN and 3 Tier-1s.	D.Foster	Done	FNAL, SARA and IN2P3.
DRC-2	31.12.05	750 MB/s data recording demonstration at CERN: Data generator >> disk >> tape sustaining 750 MB/s for one week using the CASTOR 2 mass storage system.	B.Panzer	Done	During the Christmas 2005 period the data recording rate using Castor 2 reached 950 MB. For more details see the CERN quarterly report.
Summary pf Progress					
<p>The project met the high level milestones of this quarter.</p> <p>SERVICE CHALLENGES: The major activities were the final tests for SC3, demonstrating a marked improvement in operational reliability and availability over the situation at the end of 2004, and the planning of the site infrastructure (networking, hardware, software) for the start of SC4 in June 2006. See the sites quarterly reports.</p> <p>CASTOR2: As part of the SC3 tests all experiments are now using CASTOR 2. Difficulties were experienced by each experiment during the beginning of the service phase of SC3, but these were rapidly resolved (generally within 2 weeks of the start of the experiment's operation). CASTOR 2 has also been used for the general production of some experiments, although the schedule for full migration from CASTOR 1 agreed in the summer had to be re-scheduled for the first quarter of 2006. The data recording tests at CERN at the end of the year demonstrated that CASTOR 2 is a major improvement over the previous version in terms of load balancing and error recovery, as well as showing that it is capable of achieving 50% of the nominal LHC recording requirement us Deployment at all the sites using CASTOR is planned for Q1 2006.</p> <p>LCG3D: The Distributed Database Deployment (LCG3D) project was launched, and will start with a workshop in January 2006 in order to define the goals of the project and hardware and software needed at the sites hosting the distributed database.</p>					

Outstanding Issues since Last Report					
<p>The MB decided that SC4 should have a much more detailed plan for sites, services and experiments than was done for SC3, in order to define and agree in detail and in advance which services and equipment would be available.</p> <p>A VO Boxes workshop is scheduled for January 2006 in order to clarify exactly the functionalities (permissions, network access, ports, etc.) needed by the experiments and the constraints (control, security, etc) from the sites.</p>					
Milestones Changes and Actions					
<p>OPN-2: The sites that will be on GEANT will be CNAF and FZK but a third sites will not be operational. The negotiations with GEANT are still ongoing (Feb 2006).</p> <p>SC3 rerun: In order to consolidate and confirm the results of the SC3 tests in 2005 the MB decided to repeat the tests for disk-T1 disk transfers at the nominal rate and add disk-T1 tapes at least half of the nominal rate</p>					
References and Hyperlinks					
<p>Planning information, Milestones Plans, SC4 Planning https://cern.ch/twiki/bin/view/LCG/Planning</p> <p>Management Board Minutes https://cern.ch/twiki/bin/view/LCG/MbMeetingsMinutes</p>					
Milestones for Next Quarter					
ID	Date	Milestone	Coord.	Status	Comments
SC3-3	28.02.06	Performance and throughput tests complete	J.Shiers		CERN-disk > network > Tier-1-disk and tape. Goal is to maintain for one week an average aggregate throughput of 1 GB/s from disk at CERN to disk at the Tier-1s; each Tier-1 capable of accepting 150 MB/sec to disk and 50 MB/sec to tape. All Tier-1 sites must participate. At least 5 Tier-1s must satisfy individual site throughput goals.
SC4-2	28.02.06	Use cases and service level support defined for SC4 SC workshop at CHEP	J.Shiers		Defines in detail: - the SC4 success criteria for each Tier-1(SC4-5) - the LCG Services operations (IS-1).
CAS-1	15.03.06	Castor2 Readiness Review	A.Cass		
SC3-4	31.03.06	All services on all Tier-1 sites monitored	J.Shiers		
SC3-5	31.03.06	Proposal on availability levels specified in Annex 3 of the WLCG MoU (adjusted for sites that do not provide a 24 hour service)	J.Shiers		
OPN-2	31.03.06	Tier-0/1 high-performance network operational at CERN and 6 Tier-1s, at least 3 via GEANT.	D.Foster		FNAL, SARA, IN2P3, TRIUMF, BNL and CNAF. CNAF and FZK on GEANT.
Comments and Additional Information					

QUARTERLY STATUS REPORT									
Project Name					Date				
ASGC					2-Feb-2006				
Report Period					Author Name				
2005 Q4					Jason Shih				
Milestones for the Quarter									
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments	
		138	23	40	24			network capacity at: 2Gbps	
	15.10.05	VOBoxes Installed					complete		
	17.10.05	LFC Upgraded to 2.6.0					complete		
	09.11.05	160 CPU start install					complete		
	09.11.05	Start procurement of disk extension					complete		
	20.11.05	FTS installed					complete		FTS 1.3 with Oracle backend
	30.11.05	160 CPU Made Operational					complete		
	30.11.05	298	23	40	24		complete		
	15.12.05	Start T2 SC testing (network & disk-disk)					complete		Due to internal H/W problem, ASGC continue rerun the WAN disk transfer during Feb, and continue providing sustained throughput at 110+ MB/s
	01.10.06			100					
Summary of Progress									
<p>With new procurement, we've expanded STK, disk pool of CASTOR MSS (ver. 1), to 36T and will start operation at Feb 3. Besides to this, computing capacity also increasing upto 420KSI2k after 80 dual Quanta computing nodes delivered. We plan to expand Tape system to 280T mid of Mar., and start data migration before Apr., start of SC4.</p> <p>Due to the delay of certifications from IBM, the procurement of MSS expansion have been postponed as this late, and 4 new tape drives (LTO3) and 700 LTO3 cartridges will be implemented after checkpoints setted, end of Feb. Moreover, all baseline services (VOBOX, FTS and LFC etc.) have been deployed before SC3 start, for CMS SC, we've reach sustained data transfer rate at 80+MB/s based on CASTOR+SRM, for initial phase of Atlas DM (based on VOBOX), the data transfer rate reduced to 25MB/s only.</p> <p>During SC rerun (from 1/30 and continue to Feb 11), we've noticed that the average throughput have ramped up to 90+, with peak over 100MB/s.</p> <p>To have best performance of data management over disk servers and MSS, extra I/O tuning are considered as well as TCP performance testing, several kernel parameters and SRM configuration have been introduced to improve the throughput.</p>									
Outstanding Issues since Last Report									
<p>Open issues: during SC3 rerun, we've found that the throughput is restricted to 100MB/s, however, with specific kernel parameters applied on SC storage resources, we can reach over 1G/s with iperf connected from CERN OpenLab. We're tracking down this problem and several possible solution have been proposed to resolved this situations. The validation will continue mid of Feb.</p> <p>During Atlas SC, where DM rely on VOBOX mainly to CASTOR-SC resources, we've noticed the throughput reduce to 25 MB/s only, could this be the defects of Atlas DQ? as we already reach the requiremend setup before in CMS SC with same infrastructure.</p>									
Milestones Changes and Actions									
<p>WAN to tape rerun will postpone after new MSS expansion completed.</p> <p>Schedule of B/W upgrade and integration changed.</p> <p>CASTOR-2 testing start at mid-Jan and finish deployment end of Feb., and in progress evaluation of database backend have been start in the mean time.</p>									

References and Hyperlinks									
Milestones for Next Quarter									
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments	
	01.10.06			100					
	15.01.06	Start installation of new core router and migration of local loop							
	30.01.06	Complete T2 SC testing and start peak transfer tests					In progress	Only IHEP have done initial throughput testing, and opening channels for other T2s	
	30.01.06	Complete 16TB disk procurement					complete		
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)					complete	SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA	
	15.02.06	16 TB Disk Operational					In progress	Start operation at 1/27, 2006	
	15.02.06	298	39	70	24		complete	network capacity at: 2Gbps	
	28.02.06	Complete installation of new core router and migration of local loop							
	01.04.06	298	45	100	280	75			
	31.10.06	840	245	100	280	75			
	31.12.06	950	400	100	280	75			
Comments and Additional Information									

QUARTERLY STATUS REPORT								
Project Name					Date			
CC-IN2P3					01-Mar-2006			
Report Period					Author Name			
2005 Q4					Fabio Hernandez			
Milestones for the Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
IN-05-1	1.10.05	Begin evaluation process for new compute nodes and disk servers					Ongoing	
IN-05-2	1.10.05	VO Boxes installed and delivered to experiments					Complete	As requested by Alice, Atlas and CMS
IN-05-3	17.10.05	LFC installed and tested					Complete	Service published in the information system
IN-05-4	31.10.05	FTS installed and tested by experiments					Complete	Close interaction with experiment representatives for creating channels
IN-05-5	31.12.05	410	87		200		Complete	
IN-05-6	31.12.05	Evaluation process for compute nodes and disk servers completed					Ongoing	Due to some difficulties on getting the appropriate hardware from the providers, the evaluation process for compute nodes will be completed by end of January. The evaluation for disk servers is starting and should be finished by end of February.
IN-05-7	31.12.05	Dedicated network link to CERN of 10 Gbps in service					Complete	The link was not delivered as planned. It was delivered on 17/01/2006.
Summary of Progress								
Status of the project at the end of 2005:								
Total Site Capacity: 1400 kSI2000, 300 TB disk, 1230 TB MSS, 1 Gbps shared network link to CERN								
Fraction of resources allocated to LHC experiments: 410 kSI2000, 87 TB disk, 200 TB MSS.								
The worker nodes are shared for local and grid users, and for LHC and other experiments. They are running SL 3.0.4 and LCG 2.6.0 centrally installed on AFS, so a single image of the middleware is shared by all worker nodes. In addition the site contributes to the gLite pre-production system.								
All the software required by the experiments have been installed. It includes: SRM/dCache, LFC, FTS, BDII, CE, R-GMA, etc. Started deployment of specific xrootd service as requested by Alice. In addition, 3 VO boxes have been installed and delivered to Alice, Atlas and CMS.								
Participation in SC3 with SRM/dCache, FTS and PHEDEX for CMS. Besides, the site joined the data base replication testbed of the LCG 3D project using the newly deployed Oracle 10g software.								
Organisation of a workshop on 14,15/12/2005 for coordinating the efforts of all sites contributing or planning to contribute resources for LCG in France (1 Tier-1, 3 Tier-2, 2 Tier-3).								
Outstanding Issues since Last Report								
The main issue for 2005Q4 was the unavailability of the dedicated network link Lyon-CERN at 10 Gbps. It was expected to be operational by the end of 2005 in order to be used in the SC3 re-run in week 3 of 2006. Fortunately, it was finally delivered and put in production on 17 January 2006, just in time for SC3 rerun.								

Milestones Changes and Actions								
References and Hyperlinks								
Milestones for Next Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
2006								
	01.01.06			150		na		
IN-06-1	15.01.06	Begin procurement of additional tape drives and cartridges					Complete	Although initially planned for 2006Q1, this procurement was performed in advance in December 2005 for budgetary reasons.
IN-06-2	15.01.06	Begin evaluation for automated cartridge library/tape drives upgrade						
IN-06-3	17.01.06	Dedicated network link Lyon-CERN (10 Gbps) operational						
IN-06-4	31.01.06	End of the evaluation process for compute nodes						
IN-06-5	15.02.06	Begin formal procurement process for compute nodes						
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)				SRM 2.1, VOBOXES, LFC, FTS, CE, RB, BDII, RGMA		
IN-06-6	28.02.06	Additional tape drives installed and operational						This task should be complete before the deadline
IN-06-7	28.02.06	End of the evaluation process for disk storage hardware						
IN-06-8	15.03.06	Begin formal procurement process for disk storage						
Comments and Additional Information								

QUARTERLY STATUS REPORT				
Project Name			Date	
CERN			13-Jan-2006	
Report Period			Author Name	
2005 Q4			Bernd Panzer-Steindel	
Milestones for the Quarter			Status	Comments
DR-1	31.12.05	Definition of T0 building blocks The building blocks of the Tier-0 Centre are defined, their performance and their interactions are established (disk server, tape server, tape drives and the Castor software), decision on the best way to solve the 'impedance' problem between disk and tape server.	delayed until the 15th of February	
DR-2	31.12.05	750 MB/s data recording demonstration at CERN Data generator □ disk □ tape sustaining 750 MB/s for one week using the CASTOR 2 mass storage system. Internal milestone for DRC2. Use the building blocks to setup a system to run at 750 MB/s to tape for at least a week, where 750 MB/s means the average over one week with a maximum of 4 periods of a maximum of 12 h (each) where the speed drops below 700 MB/s (minimum 600 MB/s).	achieved with 950 MB/s for one week	
DR-3	31.12.05	T0 buffer performance of 500 MB/s Expand this system to the T0 buffer setup. The disk pool would be filled by about 100 streams and in parallel the data would be read by three different client systems (emulation of the T1 export, tape writing via Castor, emulation of reconstruction).	delayed until 15th February 2006	This infrastructure should be large enough to cope safely with 500 MB/s (== 500 MB/s in to the pool and 1.5 GB/s out of the pool) and should also run for about one week.
DR-4	31.12.05	Production of 3 experiments migrated to CASTOR 2 The test and migration plan agreed in June 2004 is behind schedule. This milestone needs commitment from three experiments to staff their test and migration process.	delayed	
CC-1	31.10.05	Electrical Installation, including sub-station, commissioned	achieved 06.01.06	Safety verifications are causing delays.
DBS-1	30.11.05	Design for the Database Services at Tier-1 Centres	achieved	
DBS-2	30.11.05	Implementation and Testing Plan	achieved	
Summary of Progress				
<p>The major achievement was the Data Recording Challenge with Castor2 which run over Christmas. With 46 disk servers and 18 tape drives (3 different types, STK 9940B, LTO-3, IBM 3592B) a continuous data rate of 950 MB/s was achieved for a period of one week. The disk pool of 230 TB was filled by up to 80 input streams and the data was than migrated to tape, thus the load on the disk servers was 1.8 GBytes/s aggregate. After 7 days another 48 h was used for further tests where input data rate of 1.2 GBytes/s were achieved.</p> <p>Due to the late full availability of the new network infrastructure and Castor2 infrastructure the T0 test is delayed. As ATLAS was planning at the end of the year to do a very similar test, the IT and ATLAS tests is currently combined and first test are Running right now. The goal is to reach the required ATLAS and IT data rates by the first week of February.</p>				

The definition of T0 blocks was also delayed and the analysis of the conducted tests is still ongoing right now, a first report is being prepared. The status report will be published before the 15th of February.
 The commissioning of the sub-station has been done but a mistake caused unfortunately an electrical shutdown of the computer center in the beginning of January.
 The data base planning and status has been presented to the MB on the 6th of December 2006.

Outstanding Issues since Last Report

The Castor2 deployment plan is still behind schedule despite the very good progress towards the end of 2005. All 4 experiments have already some part of their production running on Castor2 as part of the SC3 setup. The disentanglement of the experiments have started so that each one has its own Castor2 instance. Due to the delays the movement of production and general users has now been merged. The milestone itself needs probably a refinement.

Milestones Changes and Actions

The 1 GBytes/s data recording demonstration has essentially already been achieved over Christmas 2005. The installation of Castor2 for some of the outside institutes has already been started.
 The delivery of the STK tape equipment is very late and the milestones TAPE-1 need to be shifted to the beginning of March.

References and Hyperlinks

Milestones for Next Quarter

ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
DTT-1	31.01.06	Architecture and Plan for the DAQ – Tier-0 Integration and Testing Document providing: 1. a detailed architecture for the integration of the four DAQ systems with the Tier-0 facility in the Computer Centre 2. implementation plan 3. testing plan with milestones to demonstrate nominal LHC data rates by end December 2006 and full operational capability by end April 2007.						
DTT-2	31.01.06	Testing Plan for the end-to-end DAQ – Tier-0 – Tier-1 system 1. testing plan with milestones to demonstrate full data path from DAQ to Tier-0, recording on tape, reconstruction, distribution to Tier-1s, recording on tape at Tier-1s 2. by end July 2006 – at least 2 experiments each with at least 3 Tier-1s (Level 1 milestone), 200 MB/sec aggregate throughput (should include the conditions database) 3. by end Feb 2007 – all Tier-1s, full functionality and nominal data rates (see Figure 3). By this time the conditions database must be included						
TAPE-1	31.01.06	Installation of SC4 tape equipment from two vendors The two pilot installations from STK and IBM with a capacity of >=5000 robot slots (each) and 40 tape drives (each). Equipment installed and commissioned.						
TAPE-1	28.02.06	SC4 tape equipment in operation. The two pilot installations from STK and IBM with a capacity of >=5000 robot slots (each) and 40 tape drives (each).						
ACQ-1	15.02.06	1600+700	150+360		1550+0			

Comments and Additional Information

QUARTERLY STATUS REPORT								
Project Name					Date			
FZK / GridKa					13-Jan-2006			
Report Period					Author Name			
2005 Q4					Holger Marten			
Milestones for the Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
	31.10.05	Procurement for expansion of Tape IO with nominal 320 MB/s started					done	4 drives LTO3 delivered.
	31.10.05	Procurement for expansion of Tape storage with 250 TB capacity started					done	Tape media delivered.
		Procurement for expansion of compute power with 500 kSI2k started					done	EU-wide call for tender running.
	30.11.05	Start disk procurement for capacity expansion of 200 TB for dCache					done	EU-wide call for tender running.
Summary of Progress								
Status of the project at the end of 2005:								
<u>Total capacities:</u>								
1550 CPUs (1.44 MSI2k)								
310 TB disk								
500 TB tape								
10 Gbps WAN/Intenet								
<u>LHC fraction:</u>								
530 kSI2k CPU								
83 TB disk								
146 TB tape								
All CPUs installed with SL 3.0.4 and LCG-2-6-0.								
Additional pre-production system with gLite.								
10 Gbps WAN / Internet (formerly for 'tests') reinstalled for production via GEANT/DFN.								
GridKa operates VO-Box like servers for each of the LHC VOs (called experiment specific front-ends in the GridKa design). Two of them (ALICE, ATLAS) have been updated with respective LCG VO-services, the detailed process for CMS and LHCb is under discussion.								
Participation in SC3 with SRM/dCache, FTS and PhEEx (for CMS).								
Successful transfer tests with Tier-2s at DESY (Hamburg), GSI (Darmstadt), FZU (Prague) & several Russian Tier-2 sites.								
Tier 1/2 meeting with representatives of German & international Tier-2s (that wish to use GridKa as Tier-1) in October 2005.								
<u>Planned capacities for LHC in 2006:</u>								
1030 kSI2k CPU								
280 TB disk								
393 TB tape								
10 Gbps WAN/Intenet								
additional 10 Gbps for T0-T1 via OPN								
Outstanding Issues since Last Report								
Milestone on 31.1.06 "Network: Lightpath to CERN" to be shifted to end of March'06. This is already taken into account as OPN-2 for Karlsruhe.								
Milestones Changes and Actions								

References and Hyperlinks								
Milestones for Next Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	WAN=>Disl (MB/sec)	Tape (TB)	WAN=>Tap (MB/sec)	Status	Comments
	01.01.06			150				
	31.01.06	Network: Lightpath to CERN						Waiting for GEANT2 progress
	31.01.06	Start procurement for expansion of Tape IO with nominal 640 MB/s						
	31.01.06	Start procurement for expansion of Write Pool with 100 MB/s						
	31.01.06				375	100		Tape: Expansion with 250 TB of existing 145 TB Tape IO: Expansion with nominal 320 MB/s of existing 300 MB/s
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM 2.1, VOBOXES, LFC, FTS, CE, RB, BDII, RGMA
Comments and Additional Information								

QUARTERLY STATUS REPORT								
Project Name					Date			
INFN					25-Jan-2006			
Report Period					Author Name			
2005 Q4					Tiziana Ferrari, Luca dell'Agnello			
Milestones for the Quarter								
ID	Date		Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
	31.10.05	2 x 1 Gethernet links between CNAF and GARR dedicated to SC traffic to/from CERN					ongoing	23-11-2005: CERN - CNAF connection upgraded to 10 Gb/s, uplink CNAF - CERN still at 2 Gb/s, to be upgraded to 10 Gb/s
	31.10.05	SRM/Castor and FTS					done	
	31.10.05	Farm middleware: LCG 2.6					done	
	31.10.05	min 1200 max 1500	50 TB Castor front-end	125 MB/s	200 TB (4 9940B + 6 LTO2	~ 50 MB/s	done	Tape 200 TB=(4x9940B + 6xLTO2 drives) CPU, the farm is shared
	15.12.05	min 1300 max 1600					done	Addition of approx. 50 new worker nodes
	31.12.05	1. upgrade to 10 Gb/s of the entire network path connecting CERN and CNAF completed; 2. Testing of the new path completed					ongoing	schedule shifted to mid February due to hardware and routing problems
	31.12.05	addition of new 160 TB of disk space, to be put in production according to the requirements of the experiments					done	
Summary of Progress								
<p>Total capacities:</p> <p>2500 CPU slots (1500 physical CPUs)</p> <p>disk: 400 TB, including the Castor buffer space</p> <p>tape: 200 TB (4 9940B + 6 LTO2 drives)</p> <p>2 Gb/s available bandwidth for LHC (CERN - CNAF)</p> <p>LHC fraction:</p> <p>up to 2500 CPU slots (1500 physical CPUs) - all worker nodes are shared</p> <p>disk: 112 TB, including the Castor front-end</p> <p>tape: 160 TB (4 9940B + 6 LTO2 drives)</p> <p>2 Gb/s available bandwidth for LHC (CERN - CNAF)</p> <p>All CPU installed with SLC 3.0.4 and LCG 2.6.0</p> <p>Additional test farm available</p>								
Outstanding Issues since Last Report								
Procurement process of new resources ongoing.								
Milestones Changes and Actions								
<p>1. Milestone for procurement of tape drives added</p> <p>2. Change of schedule for bandwidth upgrade</p>								

References and Hyperlinks									
Milestones for Next Quarter									
ID	Date	CPU (kSI2K)	Disk (TB)	VAN=>Dis (MB/sec)	Tape (TB)	VAN=>Tap (MB/sec)	Status	Comments	
	16.01.06			150 MB/s					
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)							SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA
	01.02.06	Possibile upgrade to CASTOR v2							Installation of test layout started on Dec 1 2005, testing phase startin in the 1st half of Dec
	01.02.06	new tapes purchased: 120 TB							
	01.02.06		TBD	150 MB/s	320 TB				
	01.02.06	Revision of local and national connectivity							<ul style="list-style-type: none"> • Usage of policy routing at the GARR access point • Type of connectivity to INFN Tier-2 under discussion • Backup link Tier-1 <input type="checkbox"/> <input type="checkbox"/> Tier-1 (Karlsruhe) under discussion
	15.02.06	upgrade of CERN-CNAF dedicated link to 10 Gbps							
	01.03.06	Evaluation of dCache and StoRM (for disk-only SRMs)							
	01.05.06	new drives for STK library							
Comments and Additional Information									

QUARTERLY STATUS REPORT								
Project Name					Date			
NDGF					2-Feb-2006			
Report Period					Author Name			
2005 Q4					Michael Grønager			
Milestones for the Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
	01/10/05	361	26	150	85	-		No tape test so far. See comments...
LCGARC-2	31/10/05	LCG-ARC Interoperability meeting at CERN					Complete	
T1COORD-1	29/11/05	Coordination meeting in Copenhagen					Complete	
HWNBI-1	30/11/05	NBI: 472 Core AMD Opteron cluster delivered, 25% for HEP/Tier1					complete	
HWNBI-2	30/11/05	NBI: HSM system delivered					complete	40TBDisk, +60TB Tape, to be expanded during 2006, 2007, 2008...
HWNBI-3	31/12/05	Servers for NBI dCache solution acquired					complete	
	31/12/05	361	26	150	85	-		No tape test so far. See comments...
Summary of Progress								
<p>Participants in the SC3 rerun will be: NBI, NSC and PDC using the same distributed DPM setup from SC3. dCache will be used for SC4 and a planning group has been established. Also a network planning group has been established. Hardware status by the end of the year is 450kSI2k, 44TB disk and 225TB tape. Some of it will not be deployed until 2006Q1. The Nordic Tier-1 will primarily run ARC. An interoperability task-force has been formed and ARC resources can now be exposed in the BDII. Plan is to integrate ARC-CE submission into the gLite resource broker. A VO-BOX for ALICE running on ARC sites is in pre-production in Bergen and ATLAS does not need a VO-BOX running on ARC-CEs, as needed services is integrated in the CE.</p>								
Outstanding Issues since Last Report								
Milestones Changes and Actions								
References and Hyperlinks								
www.ndgf.org -> projects -> sc3/sc4 give info and pointers to network monitoring the Nordic site								
Milestones for Next Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
	01/01/06	361	26	150	85	-		No tape test so far. See comments...
SC3	16/01/06	NSC, PDC and NBI will participate in the SC3 rerun using the successful DPM setup from summer 2005						
LCGARC-3	18/01/06	ARC-LCG interoperability meeting in Uppsala						
T1COORD-2	18/01/06	Coordination meeting in Uppsala.						
HWNBI-4	15/03/06	NBI compute and storage system in production						
NETWORK-	15/02/06	Network coordination meeting						

SC4-1	15/02/06	Initiate dCache test setup						
LCGARC-3	23/03/06	gLite RB patch for LCG->ARC demonstrated						
LCGARC-4	24/03/06	ARC-LCG interoperability task-force meeting at CERN						
DCACHE-1	27/03/06	dCache Single DB, distributed Pools installed						
VOBOX-2	31/03/06	ALICE-ARC VOBOX installed also at NBI						
NETWORK-2	31/03/06	Upgrade to use a single address space for Nordic sites, upgrade to use dedicated connections from sites to geant – completed						
	31/03/06	505	52	150	225	75		
SC4-2	03/04/06	SC4 T0-T1 disk-disk transfer test begins						
SC4-3	24/04/06	SC4 T0-T1 disk-tape transfer test begins						
GLITE-1	01/05/06	Upgrade of services to gLite3 begins						
DCACHE-2	30/05/06	Distributed dCache setup ready						
Comments and Additional Information								
<p>The ARC-LCG interoperability group has a wiki – see: https://uimon.cern.ch/twiki/bin/view/LCG/LCG-ARC for minutes etc.</p> <p>The dCache setup for the April SC4 test will be a quite simple and preliminary setup consisting of a central db-server and pools distributed over the participating sites. The setup aimed for June is currently under investigation in dialog with the DESY dCache team and will be completed in May.</p> <p>The ARC-ALICE VOBOX has been developed by Csaba Anderlik at the University of Bergen where it is currently being tested for pre-production.</p> <p>The hardware status consists of hardware intended for Tier-1 use in DK, FI, NO and SE. It is installed with either LCG or ARC middleware. The storage has been installed with DPM for the SC3 rerun, but will be reinstalled with dCache for SC4 in 2006Q1. The network to tape transfer rate has not been tested yet but is estimated to be varying 30 to 70 depending on the specific sites.</p>								

QUARTERLY STATUS REPORT								
Project Name					Date			
PIC					16-Jan-2006			
Report Period					Author Name			
2005 Q4					Gonzalo Merino			
Milestones for the Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
	15.10.05	150	14.3	30	85	27		
	15.10.05	SRMCastor and LFC local atlas installed					DONE	Up and running since then. Published through the Information System.
	01.11.05	VOBox-atlas installed					DONE	Operative since then and being used by Atlas SGM.
	01.11.05	start deployment of Castor disk-cache expansion					DONE	
	15.11.05	Castor disk-cache expansion operational					DONE	
	15.11.05	150	24.5	30	85	27		
	01.12.05	FTSserver-atlas installed					PENDING	Not yet deployed. Experiment need delayed.
	01.12.05	start deployment of Castor disk-cache expansion					DONE	
	01.12.05	start deployment of 2 new tape drives 9940B					DONE	This action was delayed, due to provider delay. Started on 9-Jan-2006.
	15.12.05	Castor disk-cache expansion operational					DONE	
	15.12.05	new tape drives operational					DONE	This action was delayed, due to provider delay. Completed on 11-Jan-2006.
	15.12.05	150	41.5	30	85	30		Corrected planning error. The wan->tape rate can not be higher than the wan->disk rate.
Summary of Progress								
<p>The following are the main activities carried out during the past quarter at PIC:</p> <ul style="list-style-type: none"> -Testing of configurations for disk servers deployed in castor1 stage-pools using SC3 transfers from atlas, cms and lhcb. The explored issues cover topics like RAID configurations, TCP parameters tuning. -The deployment of two new tape drives 9940B, planned to happen in the first half of December 2005 was delayed and finally happened in the second week of January 2006. The reason for this was that the hardware was delivered later than expected by the provider (StorageTek). At the time of writing this report, the drives are in operation. 								
Outstanding Issues since Last Report								
<p>The following are the outstanding issues that arose during the past quarter at PIC:</p> <ul style="list-style-type: none"> -We observe a maximum of ~2-4MB/sec per stream in the transfers to PIC. This results in high fragmentation level on the stage-pools filesystems. We are currently investigating the origin of such a limit in the bandwidth we can reach with one stream. -It should be noted that the SRM service deployed up to now has ben just in the SE-tape "flavour" (SRM interface to Castor1 MSS). Now we understand that the experiments do need an SE-disk service as well, therefore we are planning to explore and deploy d-Cache for this during the 1st quarter 2006. 								

Milestones Changes and Actions									
<p>WAN2tape capacity milestones: We have corrected the values for these milestones, so that the rate WAN2tape never exceeds that in WAN2disk. Before the upgrade of the WAN link from 1Gbps to 2Gbps (Feb-2006) the WAN2tape rate has been set to 30MB/sec (ie, dominated by the WAN capacity), and after this upgrade it has been set to 55MB/sec.</p> <p>The deployment of additional CPU planned for the first quarter of 2006 has been delayed. In terms of CPU, PIC already fulfills the requirements for 2006 (150ksi2k, after the proposed modification of the MoU figures). The tape capacity milestone at 1-Apr-2006 has been modified (set now to 90TB, equal to that in the proposed modification of the MoU figures).</p> <p>Concerning the disk capacity, we plan to fill the gap with the MoU figure for 2006 (80TB after the proposed modification) releasing a tender during 1st quarter 2006. Having into account the 1 month delay between the start tender, order placing, hardware reception and hardware in production steps, we currently plan to have the expansion in production by May-2006.</p>									
References and Hyperlinks									
Milestones for Next Quarter									
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments	
	01.01.06	start deployment of FTS server with MySQL backend							
	15.01.06	start deployment of Castor2 test service							
	15.01.06	New IBM tape library delivered							
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)							SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA
	01.02.06	FTS server (atlas and lhcb) deployed							
	01.02.06	start testing d-Cache for deploying an SRM-disk							
	01.02.06	start tender for disk expansion							
	01.02.06	start deployment of additional 1Gbps WAN							
	01.03.06	new 2Gbps WAN infrastructure operational							
	01.03.06	150	41.5	60	85	55			
	01.03.06	place order for disk expansion							
	01.03.06	1st version of SRM-dCache available for testing							
	07.03.06	Castor2/SRM2.1 service ready for testing							
	15.03.06	new IBM tape library integrated within Castor							
	15.03.06	start deployment of voboxes for cms, lhcb							
	01.04.06	SRM-dCache deployed in production							
	01.04.06	start disk expansion deployment							
	01.04.06	Start of SC4-setup phase							
	01.04.06	Add tape cartridge capacity							
	01.04.06	150	41.5	60	90	55			
	01.04.06	voboxes for cms, lhcb in production							
	01.05.06	disk expansion deployed							
	01.05.06	150	80	60	90	55			
Comments and Additional Information									

QUARTERLY STATUS REPORT								
Project Name					Date			
RAL					20-Jan-2006			
Report Period					Author Name			
2005 Q4					A.Sansum			
Milestones for the Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
RAL-1	15.09.05	493	120	100	229	75		WAN 2x1 GB/s
RAL-2	15.10.05	Closed Tender for New Tape Robot					Complete	
RAL-3	31.10.05	VO boxes, LFC, FTS, CE, RB, BDII and RGMA deployed					Complete	No upgrades outstanding
RAL-4	15.11.05	Commenced Tender for January 2006/April 2006					Complete	
RAL-5	15.11.05	Commence Service Hardening Project					Started	
RAL-6	15.11.05	CASTOR2 Hardware Tender Released					Complete	
RAL-7	30.11.05	Release tender for 10Gbit upgrade to Site Edge Router					Complete	
RAL-8	30.11.05	Place order for tape robot					Complete	
RAL-9	15.12.05	CASTOR2 hardware Infrastructure Ordered					Complete	
RAL-10	15.12.05	Tier-1 10Gbit uplink switch ordered					Started	Order just about to go out door
RAL-11	15.12.05	Delivery of 3D Production Hardware					Started	Ordered - delivery expected Feb 2006. Order was delayed
RAL-12	15.12.05	Detailed deployment plan for CASTOR2 agreed					Complete	
RAL-13	15.12.05	Delivery of Resilient Hardware for Critical Services					Complete	10 servers upgraded to dual disk. Remaining funding for resilient hardware in FY05 committed to resilient 3D infrastructure
RAL-14	31.12.05	New Tape Robot Delivered					Complete	
RAL-15	31.12.05	20% of Tier-2 GRIDPP sites tested at 300-500Mb/s for 1TB bi-directional from Tier-1					Complete	3 sites were tested, but performance issues prevented all but RAL-Manchester reaching target rate
Summary of Progress								
<p>Procurements are underway for 130TB of disk and 266kSI2K of CPU and a new SL8500 tape robot has been installed. T10K tape drives for the robot have been ordered but are delayed by STK. A 4 node cluster with SAN storage has been ordered for the 3D project, but delivery is scheduled for February – later than planned. RAL has selected CASTOR as the control system to operate the tape robot and has joined the CASTOR project. It is planned that CASTOR will be available for the July service challenge throughput test (as a trial service) and full production scheduled for September – this will be a challenge for the team. RAL has been engaged in a GRIDPP service challenge testing its downstream Tier-2 sites, good progress is being made with considerable work being carried out at the Tier-2s to improve throughput.</p>								
Outstanding Issues since Last Report								
<p>Ordering and consequently delivery of the 3D hardware was delayed as the capacity orders took precedence and delays in delivery also occurred. 3D will get priority for installation in February and we are still working towards a March service. Issues were identified with transfers between dCache and DPM. These were reported to the GDB in January and were resolved a few days later.</p>								
Milestones Changes and Actions								
<p>Work on capacity disk and CPU tenders delayed the 3D hardware procurement. The order has been placed, but delivery is not expected until early February, consequently delaying 3D deployment. We still expect to have the 3D service moved to the new RAC infrastructure before the start of SC4. Procurement of the site edge router (outside the responsibility of the Tier-1 service) is running about 1 month late, however this is not expected to have downstream implications for SC4. Edge router needs to be at 10Gb in time for the 10Gb SJ5 link to CERN. Tier-1-Tier-2 tests are being impacted by poor performance between dCache and DPM. This appears to be a generic problem, not RAL specific and has been reported to the GDB [1] a fix was provided shortly afterwards and has been deployed in the UK. We do not expect to receive the T10K tape drives in Q1, because of production delays, however we continue to plan to have the T10K drives operational for July SC4.</p>								

References and Hyperlinks								
[1] http://agenda.cern.ch/askArchive.php?base=agenda&categ=a057701&id=a057701s1t17/transparencies								
Milestones for Next Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
RAL-16	01.01.06			150		na		
RAL-17	15.01.06	CASTOR2 Hardware Delivered					Started	Delivery scheduled
RAL-18	15.01.06	Delivery of Site Edge Router 10Gbit Upgrade						Tender evaluation started - delivery now scheduled for February
RAL-19	15.01.06	4*6TB disk resource allocated to dteam in preparation for SC3 throughput test						
RAL-20	15.01.06	On-Call System in Place						
RAL-21	15.01.06	Airconditioning Capacity Upgrade Installed						Now scheduled for end February. This is not a
RAL-22	15.01.06	dCache Upgraded to version 1.6.6 (SRM 1)						Follow two parallel tracks. dCache will continue to provide production SRM until at least March 2007 but at some point (possibly end 2006) will become read only.
RAL-23	15.01.06	FTS Upgraded to latest release in order to support srmcp						Depends on timely FTS release
RAL-24	15.01.06	dCache ready for 150MB/s disk to disk.						
RAL-25	15.01.06	Tier-1 10Gbit uplink switch delivered						
RAL-26	15.01.06	493	120	150	229	75		
RAL-27	31.01.06	GRIDPP sites tested at 300-500Mb/s for 1TB bi-directional						
RAL-28	31.01.06	Migration complete from old to new Robot						Functionality only
RAL-29	31.01.06	Test CASTOR2 System Running						Depends on effort from CERN
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA
RAL-30	15.02.06	dCache Upgraded to support SRM 2.1. (dCache continues to back end into RAL ADS tape system)						Depends on receiving dCache SRM 2.1 on 15th January and deploying 2-3 weeks later
RAL-31	15.02.06	3D Service Moves to Production Hardware						May be a couple of weeks late
RAL-32	15.02.06	Tier-1 Connected to Site Edge Router at 10Gb/s						
RAL-33	15.02.06	1st Disk and CPU delivery						Expected 2 weeks later
RAL-34	28.02.06	New Tape drives Installed						Not funded by Tier-1 but available for test
RAL-35	28.02.06	UKLIGHT Provisioned at 4*1Gbit						May be earlier but scheduled January date now impacted by throughput test.
RAL-36	28.02.06	Order Tape Drives and Media for new Robot						
RAL-37	15.03.06	Nagios monitoring System deployed (replaces SURE)						Better automation and escalation
RAL-38	31.03.06	Completion of Phase I Service Hardening						
RAL-39	31.03.06	Place order for second delivery of disk/CPU						
RAL-40	31.03.06	100% of Tier-2 GRIDPP sites tested at 300-500Mb/s for 1TB bi-directional from Tier-1						
RAL-41	31.03.06	Local (Preliminary) Throughput Test of CASTOR2 Complete.						Limited throughput test - 50MB/s?
RAL-42	31.03.06	ATLAS, ALICE and CMS VOs up to date.						
Comments and Additional Information								

QUARTERLY STATUS REPORT								
Project Name					Date			
SARA-NIKHEF					1-Feb-2006			
Report Period					Author Name			
2005 Q4					Kors Bos			
Milestones for the Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
	31.10.05	300	4	60	50	45		2 Rembrandt nodes with 2 TB each
	31.11.05	data migrated from teras SE to new disk space					done	the now free 10 TB disk is added to SAN
	31.12.05	purchase and installation of 5 new storage servers with 2 TB disk each					done	to replace Rembrandt nodes
Summary of Progress								
Successful participation in SC3 although with frequent network problems between Amsterdam and Geneva.								
Outstanding Issues since Last Report								
Network problem fixed by Global Crossing. Still to be seen in the weeks to come.								
Milestones Changes and Actions								
References and Hyperlinks								
Presentation by Ron Trompert (SARA) in the January GDB: http://agenda.cern.ch/askArchive.php?base=agenda&categ=a057701&id=a057701s1t7%2Fmoreinfo%2Fgdb.pdf								
Milestones for Next Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
	01.01.06	300	10	150	50	45		5 data servers with 2TB each
	06.01.06	Upgrade dCache to 166-3					done	
	15.01.06	dedicated 10 GE I/F on 10 GE link to CERN					done	
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA
	31.02.06	setup T2 tests						
	31.02.06	separate T1 tape storage from general tape storage						
	31.02.06	purchase 400 9940B cartridges						if needed
	31.02.06	purchase 4 STK9940B tape drives						total of 6 tape drives for Tier-1 only
	31.02.06	purchase general purpose servers						

	31.03.06	install central DB node for FTS and local LFC						
	31.03.06	Migrate data from old teras.sara.nl SE to srm.grid.sara.nl						
	31.03.06	new SARA 10G internal network infrastructure						all traffic through new Cisco switch
	31.03.06	300	10	150	50	90		now with 6 T1 dedicated tape drives
Comments and Additional Information								

QUARTERLY STATUS REPORT								
Project Name						Date		
TRIUMF						30-Jan-2006		
Report Period						Author Name		
2005 Q4						Reda Tafirout		
Milestones for the Quarter						Status	Comments	
		18	6.5	75	8.5	36		
	4.10.05	10G Foundry Rx-4 switch ordered					received	10 Gb/s backbone for the Tier-1 centre
	07.10.05	FTS upgrade					completed	version 1.3, with T1-T2 channels defined
	28.10.05	LFC deployment					completed	
	28.10.05	VOBOX deployment					completed	along with ATLAS DDM/DQ2
	30.11.05	DWDM optics for 10G TRIUMF/BCNET					delayed	expected delivery mi-February
	07.11.05	Hardware evaluation/testing : R200 storage unit 16 TB (NetApp)					completed	unit returned on 27/01/06
	15.11.05	SC3 phase 2: ATLAS tests, T0-T1 exercise					delayed	happened only in December of 2005
	03.12.05	SC3 phase 2: ATLAS tests, T0-T1 exercise					completed	few days only (controlled by ATLAS)
	05.12.05	Hardware evaluation/testing: dual-core blades (Dell)					ongoing	3 months evaluation period
	14.12.05	first commissioning data sent to TRIUMF					done	using ATLAS data management tools
	31.12.05	2 Persons hired to support development and operations (lead sys admin. + grid expert)					completed	new personnel effective in Q1 of 2006
Summary of Progress								
<p>The main achievements were the deployment of a VOBOX service for ATLAS, including data management software, as well an LFC, followed by successful transfers between CERN and TRIUMF of emulated RAW, ESD and AOD data files (part of the T0-T1 distribution exercise). We also received the first commissioning calibration data sets from ATLAS. In terms of networking we acquired a Foundry Rx-4 switch, an important milestone in establishing a 10 Gig link. For future considerations in terms of hardware acquisition, we evaluated the performance of a storage unit from Net Appliances, as well as dual-core processors blades from Dell.</p>								
Outstanding Issues since Last Report								
<p>10 Gig link still not fully commissioned yet, awaiting delivery of DWDM optics and also the completion of Amsterdam-CERN link portion.</p>								
Milestones Changes and Actions								
References and Hyperlinks								
<p>Information on current setup can be found at: http://grid.triumf.ca/status/sc3.html</p>								
Milestones for Next Quarter						Status	Comments	
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
	01.01.06			50		na		
	15.01.06	SC3 re-run						

SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)		SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA
	15.02.06	delivery of DWDM optics for 10 Gig link & optical tests		
	01.03.06	Ordering of two tape drives + cartridges		
	06.03.06	dCache upgrade to 1.6.x		
	13.03.06	FTS upgrade to version 1.4		
Comments and Additional Information				

QUARTERLY STATUS REPORT								
Project Name					Date			
BNL - US ATLAS					10-Jan-2006			
Report Period					Author Name			
2005 Q4					Bruce Gibbard			
Milestones for the Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
UA-1	01.07.05	500	150	150	150	60	Complete	
UA-2	01.08.05	SRM, FTS, FTS server installed and operated					Complete	
UA-3	01.08.05	LCF, ATLAS VObox installed					Complete	
UA-4	na	As ATLAS VObox appl become available they will be installed & tested As new versions of any of above become available					Ongoing	
UA-5	01.12.05	Begin procurement of expanded LAN infrastructure					Complete	
UA-6	01.12.05	Begin procurement of new Tape subsystem					Complete	
UA-7	15.12.05	Expanded LAN infrastructure order placed					Delayed	Expected to complete by 15-Jan-06
UA-8	15.12.05	New Tape subsystem order placed					Complete	
UA-9	31.12.05	OSG 0.4.0 deployed and operational					Delayed	Expect to complete by 31-Jan-06
Summary of Progress								
<p>During the reporting period the procurement process for a new mass storage system resulted in an order for a SUN/StorageTek SL8500 tape library and 10 LTO Gen 3 tape drives with delivery scheduled for early 2006 and installation and commissioning expected to complete in time for the beginning of SC4. All planning for expansion of the Tier 1 LAN from a 5 Gb/sec bonded 1 backbone to a 20 Gb/sec backbone was completed and the ordering of required equipment begun. This too, along with corresponding upgrades of the Wide Area Network, is on schedule to be operational for SC4.</p>								
Outstanding Issues since Last Report								
<p>The major outstanding issue for the Tier 1 at BNL is continued incompatibility between FTS and the SRMCP interface to dCache. Minor slippages in milestones mentioned above are not expected to effect later milestones or required services.</p>								
Milestones Changes and Actions								
References and Hyperlinks								
Milestones for Next Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
UA-10	15.01.06	Begin installation of expanded LAN infrastructure						Expect delay to 15-Feb-06
UA-11	15.01.06	Begin installation of new Tape subsystem						Expect delay to 15-Feb-06

UA-12	15.02.06	Expanded LAN infrastructure operational						Expect delay to 1-Mar-06
UA-13	15.02.06	New Tape subsystem operational						Expect delay to 1-Mar-06
UA-14	15.02.06	500	150	200	300	200		
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA
UA-15	28.02.06	BNL Wide Area Network Upgrade (OC48 -> 2xλ) operational						
UA-16	31.03.06	bs BNL Tier 1 <-> Tier 0 dedicated connectivity operati						
OPN-2	31.03.06	Tier-0/1 high-performance network operational at CERN and 6 Tier 1's						BNL expected to be 1 of these 6
UA-17	31.03.06	Expanded facility power and cooling capacity available						
Comments and Additional Information								

QUARTERLY STATUS REPORT								
Project Name					Date			
FNAL - US CMS					17-Jan-2006			
Report Period					Author Name			
2005 Q4					Ian Fisk			
Milestones for the Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
	31.10.05	786	91 +36 resilient	750 MB/s demonstrated	229	7 drives for CMS + 8 shared	Done	
		WAN=>Disk is only limited by WAN					Done	Demonstrated in SC2
	31.12.05	Hierarchical Fair Share in Condor						Delayed to 2006 Q1 (See comment below)
	31.10.05	User Disk: 6 TB managed + 6 TB physics space					Done	
2006								
	01.01.06	1000	96+47 resilient	750MB/s	320	7 drives for CMS + 8 shared	Done	
Summary of Progress								
<p>The CMS Tier-1 center at FNAL had a productive quarter in the areas of hardware procurement, deployment, and commissioning as well as the areas of service deployment. FNAL commissioned the final worker nodes of the 2005 procurement. In 2005 280 dual AMD Opteron nodes were purchased and deployed in three staged deliveries. So far performance, power consumption, and operation experience of the nodes is good. This raises the total computing capacity of the FNAL tier-1 center to approximately 1MSI2k. Additionally there were 45TB of disk deployed within the dCache system. These are NEXSAN SATA disk enclosures linked with Fiber channel to between 2 and 4 data serving hosts. These disk components are expected to be the basis of the large 2006 disk procurement. The processing procurement is expected to be dual core dual CPU Opteron nodes.</p>								
<p>With the new procurements and the existing mass storage capability, the Tier-1 center at FNAL has met the capacity milestone for Q4 2005. There are approximately 1000kSI2k of processing, 96TB of normal dCache space and an additional 47TB of raw resilient space, the wide area transfer rate was demonstrated during SC2, and there is more than 320TB of capacity in the mass storage. The only other capacity related milestone was the deployment of User Disk. There are currently 12TB of user controlled disk space, which is being heavily used by the local user community.</p>								
<p>In addition to the procurement the Tier-1 center at FNAL upgraded the grid services to LCG 2.6. There was a reasonably delay in the migration from 2.4 due to available expertise and the desire to maintain a stable service for CMS. The migration was relatively smooth, the information and monitoring infrastructure required some effort to get properly configured.</p>								
<p>The OSG 0.4 release is expected during January. The CMS Remote Analysis Builder (CRAB) has been upgraded several times during the forth quarter for use by the local community to submit to remote sites. The OSG version of the submission tool is expected in January for a user release.</p>								
Outstanding Issues since Last Report								
Milestones Changes and Actions								
<p>The milestone to enable Condor hierarchical scheduling during the final quarter of 2005 was not met and is expected to be deployed in the first quarter of 2006. This does not change the functionality of the Tier-1 center, but it increases the operational load associated with tuning priority.</p>								

References and Hyperlinks								
Two status reports were given about the progress of the Tier-1 at FNAL during the forth quarter.								
The first was given at the LHCC review and is available at http://agenda.cern.ch/askArchive.php?base=agenda&categ=a045606&id=a045606s1t13/moreinfo								
The second talk was given at the CMS Computing meeting and can be found http://agenda.cern.ch/askArchive.php?base=agenda&categ=a056558&id=a056558s1t3/transparencies								
Milestones for Next Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	VAN=>Dis (MB/sec)	Tape (TB)	VAN=>Tap (MB/sec)	Status	Comments
	15.01.06	LCG 2.7 Upgrade (CE, SE, BDII,)						
	15.01.06	Upgrade to SRM 2						
	01.02.06	RGMA (If Available for Condor)						
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA
	01.02.06	SC4 CMS Local Catalog Functionality (LFC, GRLS, Trivial Catalog, etc)						
		Third 6509 Network Switch Acquisition						
		30Gb/s Between worker nodes and disk servers						
	01.04.06	1000	96+47 resilient	750	420	7 drives for CMS + 8 shared		
Comments and Additional Information								

QUARTERLY STATUS REPORT				
Project Name			Date	
Applications Area			16-Jan-2006	
Report Period			Author Name	
2005 Q4			Pere Mato	
Milestones for the Quarter			Status	Comments
SPI-1	31.12.05	Provide the tools for generating CMT and SCRAM configurations from a common generic configuration description based on XML description files. Be able to update the web and distribution's kits from the same description.	Partial completion	Tools were developed, to provide from a common configuration description based on XML files the following outputs: a) SCRAM and CMT configuration information ("mocom") the part for scram is already in use, the deployment for CMT is planned for Q1/06 after discussion with the maintainers of CMT and the main users. b) a web page to show the configuration of the external packages for a given configuration tag (http://spi.cern.ch/config/viewConfig.html) Work on generating the distribution kits from the same
ROOT-1	30.09.05	Make available prototypes addressing different topics for the SEAL+ROOT merge (Math libraries, Dictionary libraries, etc.) such that detailed planning for the experiments migration can be established. These prototypes should be available by the ROOT workshop at the end of September.	Completed	The libraries MathCore and MathMore, Reflex and Cintex have been migrated from the SEAL CVS repository to ROOT CVS.
ROOT-2	30.09.05	Demonstration of the new the Parallel ROOT facility (PROOF) in a cluster of 32 CPU's provided by CERN/IT. This new version of the system should include asynchronous queries, GUI session controller, interactive batch mode.	Completed	PROOF in asynchronous mode has been demonstrated at the ROOT workshop in September. The new GUI session controller is part of ROOT 5.08 released in December.
ROOT-4	31.12.05	Finalization of the fitting and minimization application programming interfaces and integration of the new C++ implementation of Minuit in the	Completed	The Minuit2 package originally in SEAL has been reworked and interfaced with the ROOT Virtual fitter. It is part of ROOT 5.08. Fitting in ROOT can be done with the previous versions of Minuit and Fumili as well as the new versions Minuit2 and
POOL-1	31.10.05	Production quality release of the relational database API (RAL) package, which should include the new interface recently reviewed.	Completed	The first public release of CORAL, was made available on 22/11/2005 (version 1.0.0) which contained all of the new functionality exposed through the reviewed interfaces. Prior to that there had been a few internal releases allowing beta-testers from the COOL development team and the experiment software integrators to send early feedback. Since the first public release, two other main and one minor (bug-fix) ones have been made available. As of CORAL version 1.1.0, a plugin based on the frontier-client is included in the releases. The plugin allows to transparently use FroNtier servers (possibly cached via SQUID) as read-only Database backed. This new functionality has been tested by the CMS experiment with the full POOL software stack (including object relational mapping) as part of their LCG 3D work and several enhancements went into the Frontier code and the first production release of the plugin in POOL 2.3.
POOL-1	31.12.05	POOL framework based on new C++ reflection libraries (Reflex) available for the experiments to be used in production. Validation by the experiments completed.	Completed	The POOL code is based on SEAL Reflex as of version 2.2.0, released on 21/9/2005. Since then 7 bug-fix releases have been produced which allowed the experiments to gradually pick up this part of the SEAL software through POOL. POOL will migrate to the ROOT version of Reflex (moved with SEAL 1.8) with the upcoming POOL 2.3 release.

COOL-1	30.11.05	Conditions Database (COOL) release based on the latest version of RAL including bulk insertion operations and extended tagging functionality.	In progress	COOL release based on CORAL will be released this week (CORAL_1_2_1 released on January 14 was the first release with full support for all three backends) Bulk insertion functionality is not yet implemented in COOL. Its implementation in COOL depended on the implementation in RAL/CORAL of bulk update/delete, which was only provided in CORAL (not in RAL). This is the first new feature available in CORAL that will need to be picked up in COOL (for providing bulk insertion). Extended tagging functionality has not been implemented yet.
COOL-2	31.12.05	First prototypes of API and command line tools for data extraction and cross-population of COOL databases. These tools are important for supporting partial or complete distribution of the experiment's conditions databases with several databases technologies.	Completed	Basic prototype of a data extraction and copy tool has been implemented in the new package PyCoolUtilities. Basic prototype of a data inspection tool has been implemented in the same package
SIMU-1	15.12.05	Apply the Fluka-Geant4 (Flugg) geometry interface to one of the LHC calorimeter test-beam simulation.	In progress	It is now technically possible to get hits from Fluka and then digitize and analyse them exactly as for Geant4. A very preliminary plot comparing the visible energy obtained with the two simulation engines (Fluka and Geant4) has been produced. Some cross-checks are now needed, together with a proper selection of the thresholds used in Fluka in order to be as close as possible to the Geant4 configuration. After that, a production with Fluka, and one with the latest Geant4 version can start. A presentation of these developments will be scheduled in the next Physics Validation meeting. Final results are expected for the first quarter of 2006.
SIMU-2	15.12.05	Production quality release of the MC generator level production framework.	Delayed	Software was recently reviewed by both LCG and CMS (Q3 2005): dependencies introduced at a preliminary stage turn out not to respect the agreed design. A common CMS/LCG interest to develop the framework in a single branch was re-verified (Q4 2005). The pursued design is confirmed. Support on the CMS side has been increased (supervision of the current PRS Generators conveners, development in the context of the new CMS-SW). Support on the LCG side has been also increased (GENSER integrator now acts as advisor for interfacing/testing each single Monte Carlo package). Next official report on the status of the software will be given to the LCG Generator meeting of January 2006. Verifications on the overall status of the package are expected at the end of each quarter.
SIMU-3	20.12.05	New Geant4 public release including positron annihilation and geometry voxelisation improvements in addition to the regular bug fixes and small improvements included on each release.	Achieved	Deliverable achieved as release 8.0 of Geant4 was made on December 16th, with key scheduled developments, including several developed by team members, or in which team members contributed. A key release feature was the significant revision of the multiple scattering model and process. Team members tested this for typical HEP calorimeter test cases. New correction factors were introduced in the ionisation process for heavy ions and exotic particles. New transition radiation models were created for straw tubes (ATLAS), and for transparent regular (ALICE) radiators. In the geometry we added an option to check for overlaps when a volume is constructed. New constructors for volumes and solids enable direct persistency with Pool/Root. Nested volume parameterisations can be used to navigate fast and use less memory use for regular or voxel-like structures. In hadronics, the precise high energy elastic scattering was extended to protons. New processes were created for capture of negative particles, and for charged current interactions of neutrinos.

				The physics lists were revised adapting to new design of particles. The EM components used in the physics lists were consolidated, and variants provided with processes of release 7.1 to enable smooth user transition. The toolkits was migrated to support CLHEP 2.0.2.2, while keeping support for 1.9.2.2, and to use <sstream> in place of <strstream> in the kernel and examples. Testing for this release, included system testing and the extended GRID regression tests for different release candidates identified several problems and issues (five in the extended regression tests), which were addressed in time for the release.
SIMU-4	31.12.05	First results of the ATLAS combined and 2004 test-beams data comparisons.	In progress	Many activities are undergoing in ATLAS to analyse the 2004 combined test-beam data, and the first results are now expected for the first quarter of 2006.
Summary of Progress				
<p>During this quarter the main work of consolidation of the AA software has continued, as well as the support for production deployments in the experiments.</p> <p>The migration to the new Reflex library for the persistency framework (POOL) and Python scripting has been completed. Several iterations of POOL have been made available to the experiments to validate the changes. The final version of Reflex and Cintex has been released as part of ROOT 5.08 in December and will be used instead of the SEAL ones in the coming weeks. At the same time the appropriate components will be removed from SEAL releases and this will be an important milestone towards the completion of the SEAL and ROOT merge.</p> <p>The first public release of the new re-engineered version of the relational database API package (CORAL) was made available. The adaptation of POOL and COOL (conditions database) to this new package is ongoing and finishing soon. This will be on time for the experiments to be integrated in their production software to be used in this year's major data challenges.</p> <p>End of the year production versions has been released of ROOT and Geant4 which include a long list of new functionality required by the experiments.</p> <p>Andreas Pfeiffer has replaced Alberto Aymar as SPI project leader. The project scope has been reassessed during the quarter. The reduction of exclusive manpower for SPI activities is being handled with the partial assignment of people for which the main involvement is in software development projects.</p>				
Outstanding Issues since Last Report				
None				
Milestones Changes and Actions				
Reschedule SIMU-1 to 31.03.06				
SIMU-2. It is proposed to be renamed as "New generator level production framework: beta release", and the expected due date: 30 June 2006.				
Reschedule SIMU-4 to 31.03.06				
Reschedule COOL-1 to 31.03.06				
New Proposed Milestones				
SPI-?	31.03.06	Provide the interconnection/interoperability between the savannah and HyperNews services.		Clarify the needs of the experiments on the interconnection/interoperability between the savannah and HyperNews services and provide an initial implementation
SPI-?	31.03.06	Generate CMT configuration and distribution kits from the common (XML based) configuration description.		Completion of the milestone SPI-1
SIMU-?	30.06.06	Investigation of correction for test-beam data for validation of stand-alone simulation engines (VD617)		Investigate the possibility to get, for at least some observables, e.g. the shower shapes, corrections of data for digitization and reconstruction effects, in such a way to allow direct comparison with stand-alone simulation, provided with the correct geometry imported from a GDML file. As a concrete case, the ATLAS electromagnetic barrel calorimeter test-beam of 2002 will be considered

SIMU-?	30.06.07	Application of corrections of test-beam data, for validation of stand-alone simulation, to the LHC calorimeter test-beams (VD703)		In the case of success for milestone SIMU-6, then generalize the machinery to apply to other test-beam analyses, providing to the detector simulation developers important validation data to be preserved through the years
ROOT-?	31.03.06	Have the rootcint dictionary code generator interfaced with the Reflex and gccxml options		The CINT data structures will be filled automatically from the Reflex data structures via the Cintex gateway. This is the first step to allow serious work in adaptation of the CINT interpreter to Reflex.
ROOT-?	30.06.06	First version of CINT running directly with the Reflex data structures as part of the ROOT June release		The time scale for this milestone is very tight. If we realized by end of April that we cannot meet this milestone, it would have to be postponed to the December release.
ROOT-?	30.06.06	The new Fit GUI released as part of the ROOT June release.		At the same time an extension of the Virtual Fitter that uses the new libraries in MathCore and MathMore will be developed.
ROOT-?	15.12.06	PROOF demonstrated in production in at least one of the LHC collaborations.		This new version should include monitoring of the resources, interfaces with a priority scheduler, dynamic allocation and deallocation of servers. An improved client/server I/O is available using read-ahead combined with the file cache.

Additional Information

Milestones for Next Quarter		Status	Comments
SPI-2	28.02.06	Provide a web based "user discussion forum" service interfaced with Savannah. This new service should allow projects and experiments to easily setup and manage discussion subjects.	
ROOT-5	31.03.06	The Python interface to ROOT (PyROOT) adapted to directly use the new C++ reflection library (Reflex). This would avoid the intermediate software layers and additional dependencies of the current implementation, improving the overall design and maintainability.	
ROOT-6	30.04.06	The ROOT C++ interpreter (CINT) adapted to use the new C++ reflection library (Reflex). Applications will require a single dictionary with reflection information in memory. Backward compatibility will need to be provided to old ROOT and POOL applications.	
COOL-3	31.03.06	POOL overall performance study and validation of the experiments requirements. This study should identify the areas that will require further work and optimization.	
SIMU-5	31.03.06	Monte Carlo event generator files database (MCDB) publicly available and able to deal with large files.	

Comments and Additional Information

QUARTERLY STATUS REPORT				
Project Name			Date	
Deployment Area			9-Feb-2006	
Report Period			Author Name	
2005 Q4			Ian Bird	
Milestones for the Quarter			Status	Comments
GD-1	31.10.05	Operations monitoring metrics for EGEE grid agreed EGEE metrics – scheduled to be agreed during the Pisa conference	Done.	A proposed set of performance and reliability metrics was agreed. Effort to work on them was identified.
GD-2	15.11.05	SC4 detailed plan agreed (following November GDB)	In progress	T0-T1 OK T1-T1 discussed in CHEP06. This was agreed at the basic level in November. Now in the process of detailed planning of middleware to be available in SC4, and services to be deployed at each site. To be agreed at CHEP.
GD-7	30.11.05	Decision on policy for interoperability between EGEE and OSG	Done.	Cross job submission (including common information publishing schema) has been demonstrated. Agreed with US ATLAS and US CMS for OSG. Clarify if this is a requirement of ATLAS and CMS. Cross operations monitoring principles have been discussed at the Abingdon joint operations workshop (Sep 05).
GD-3	31.12.05	Tier-1 progress reporting in place (November GDB)	Done.	By the PO starting to report to the MB end of Jan. 2006
EXT-1	30.10.05	(DESY) dCACHE version suitable for SC4	In progress	Will be a staged deployment as SRM v2.1 functionality becomes available. Global space reservation will not be there for SC4
EXT-2	30.11.05	(EGEE/JRA1) consolidated VOMS support in place and operational	In progress	This is in LCG-2.7.0 and will be tested in production
EXT-3	31.12.05	(EGEE/JRA1) New RB (bulk job submission) through certification process – and tested by ATLAS or CMS on preproduction service	In progress	Version in pre-production. Needs gLite1.5 - expected to be on PPS in March.
EXT-4	30.11.05	(EGEE/JRA1) New CE through certification process – and tested by at least one experiment on pre-production service	In progress	Same as above.
EXT-5	30.11.05	(RAL/CERN) CASTOR 2 SRM – basic version	In progress	Test version provided. Experiments need to test.

Summary of Progress				
<p>Status of the project, major key achievements, other work completed during the period.</p> <ul style="list-style-type: none"> - major activity was the support for SC3 production use by the experiments. - pre-production service and certification testbed: significant effort to deploy and debug the gLite 1.4.0 and 1.4.1 releases - preparation of the LCG-2.7.0 release 				
Outstanding Issues since Last Report				
Milestones Changes and Actions				
<p>Comments on the coming milestones, changes to the plan or additional milestones.</p> <ul style="list-style-type: none"> - ongoing detailed planning for SC4 and LCG pilot service; functionality expected in major middleware distribution releases anticipated in 2006 				
References and Hyperlinks				
<p>Additional pertinent information (web links) and other documents that contain useful details on the status of project.</p> <ul style="list-style-type: none"> - link to summary of experiment requirements on Baseline Services: https://uimon.cern.ch/twiki/bin/view/LCG/SummaryOpenIssuesTF - links to SC4 planning: https://uimon.cern.ch/twiki/bin/view/LCG/TaskForces 				
Milestones for Next Quarter			Status	Comments
GD-4	15.01.06	<p>Finalise the Baseline Services specification for the initial LHC service</p> <p>Includes specifying VO Boxes and deciding on the requirements for VOMS roles and groups; understanding the implementation issues and the associated development schedule; defining a delivery schedule for the</p>	In progress	The report from the Baseline Services workgroup will provide new milestones on the different services
GD-9	31.01.06	<p>OSG and NDGF Baseline Service Plans</p> <p>Plan of the baseline services on the different grid implementations for SC4</p>	In progress	This must be provided by NDGF and by US ATLAS and CMS for OSG. OSG plan discussed in OSG consortium meeting in January.
GD-10	31.01.06	<p>Operations monitoring metrics implemented</p> <p>Covering the metrics and criteria for</p>	In progress	The metrics agreed in GD-1 milestone. Implementation in SFT and Gridview
GD-8	31.01.06	<p>Baseline services available for setup phase of SC4</p>	In progress	LCG_2_7 deployment starting at end January, gLite 3.0 on PPS in March.
GD-5	28.02.06	<p>System and application tests for SC4 integrated in the Site Functional Test (SFT) framework</p>	In progress	
GD-6	31.03.06	<p>Service availability measurement system in place</p>	In progress	Depends on GD-5
EXT-6	31.01.06	<p>(RAL/CERN) CASTOR 2 SRM for SC4</p>		Experiments should clarify the granularity and the kind of access control
	31.01.06	<p>dCache SRM for SC4 ready for test by the experiments</p>		
EXT-7	TBD	<p>Implementation of VOMS roles and groups – placeholder - this can be defined only after final report of the baseline services group (end Nov 05)</p>	In progress	Implementations are understood: LCAS/LCMAPS in EGEE and GUMS etc in OSG.
Comments and Additional Information				

QUARTERLY STATUS REPORT				
Project Name		Date		
ARDA		16-Jan-2006		
Report Period		Author Name		
2005 Q4		Massimo Lamanna		
Milestones for the Quarter			Status	
			Comments	
1.6.11.1	31.12.05	ALICE: Use the gLite middleware (version 1.2) on the extended prototype (eventually the pre-production service) and provide feedback (technical issues and collect high-level comments and experience from the experiments)	delayed	We are still using previous versions of gLite/LCG. For specific testing, we have an ongoing collaboration with experiments task forces (e.g. ATLAS: WMS measurements)
1.6.11.2	31.12.05	ATLAS Use the gLite middleware (version 1.2) on the extended prototype (eventually the pre-production service) and provide feedback (technical issues and collect high-level comments and experience from the experiments)	delayed	
1.6.11.3	31.12.05	CMS: Use the gLite middleware (version 1.2) on the extended prototype (eventually the pre-production service) and provide feedback (technical issues and collect high-level comments and experience from the experiments)	delayed	
1.6.11.4	31.12.05	LHCb: Use the gLite middleware (version 1.2) on the extended prototype (eventually the pre-production service) and provide feedback (technical issues and collect high-level comments and experience from the experiments)	delayed	
Summary				
<p>The LCG ARDA group is developing four prototypes for the LHC experiments ALICE, ATLAS, CMS and LHCb. This document reports on the activity of the ARDA group in the fourth quarter of 2005. The activity in 2005 has the goal to migrate of current applications (ARDA prototypes) to gLite and to involve LHC users. An overview can be found in the SC2 presentation last October: http://lcg.web.cern.ch/LCG/activities/arda/public_docs/2005/Q4/ARDAatSC2-051014.ppt.</p>				
<p>ALICE The different components of the ALICE prototype are being integrated together. A preview has been shown during the LCG Comprehensive review, showing examples of user analysis. At the same time, a demo was run at SuperComputing 2005 and the ARDA activity on efficient access on Grid data has been presented at the SuperComputing conference (presentation and paper). We expect the user guide to be available at the beginning of February 2006.</p>				
<p>ATLAS The different components of the ATLAS prototype are being integrated together. A preview has been shown during the LCG Comprehensive review, showing examples of user analysis (via production system and Ganga). At the same time, a demo and posters have been presented in the EGEE conference in Pisa. Successful presentations and demos have been run for pilot ATLAS users. Within the ATLAS task force, ARDA has contributed middleware studies (WMS bulk submission) and some preliminary work on the new LEXOR version (together with the INFN Milano group).</p>				
<p>CMS The activity with CMS is continuing as foreseen. Since the ARDA analysis system was already shown (LHCC in May), the CMS demo was run by a CRAB user (CRAB is the official CMS system for analysis). It should be noted that the monthly number of jobs for CRAB and ASAP has grown to a remarkable 150k grid jobs/month (more or less evenly split between the two systems). We remind that the key ASAP features will be contributed in to the next CRAB version. The CMS dashboard activity is continuing, notably focusing on improving the performance for complex querying and streamlining the database back end from a PostgreSQL prototype hosted on a desktop PC to the IT Oracle system.</p>				

<p>LHCb Ganga has been demonstrated both at the EGEE conference and in the LCG review. In addition the usual activity with LHCb users has continued. On the development side, we underline testing contributions (ARDA-Russia visitors), a new PBS back end and the splitting contributed by a user. In addition, a brand new QT-based user interfaced has been developed and integrated.</p>				
<p>Metadata Activities The AMGA system is now part of gLite 1.5 candidate release. The activity during his quarter was concentrated on streamlining the testing part and providing additional functionality in particular on the security part (notably the database can be updated only presenting appropriate VOMS certificates while the reading is possible with no authentication to maximise performances).</p>				
Outstanding Issues since Last Report				
<p>Preparation of the activity of the section IT/PSS/ED (containing the ARDA persons) being done. ATLAS and CMS met so far. Realistically a new programme of activities will be agreed in the first weeks of February</p>				
Milestones Changes and Actions				
<p>Milestones and personnel effort The delayed milestone (1.6.10.4: "LHCb: Use the gLite middleware (version 1.1) on the extended prototype (eventually the pre-production service) and provide feedback (technical issues and collect high-level comments and experience from the experiments)") has been reached by exposing LHCb collaborators to our activity on the WMS (activity continuing independently in LHCb). The milestones 1.6.11.1/4 (usage of a new version of gLite) could not be reached as such (the gLite version 1.5 is becoming available in the middle of January. On the other hand, the demonstrations of the different prototypes (EGEE conference and LHC review) show the evolution of the prototype and some coherence with the production systems (e.g. GANGA/DIRAC, CMS dashboard aggregating data from users jobs, production system and SC3 jobs). As foreseen in the last quarterly report, the definition of the activity for 2006/7 has started and it should be finalized by February 2006. On the personnel side, we acknowledge the important contribution from ASCC Taiwan (now ASGC). ASGC has committed resources to contribute mainly in the ATLAS and CMS prototypes; the two teams will provide about 2 FT The contribution from the Russian institutes was featured by several very successful visits during this quarter.</p>				
References and Hyperlinks				
<p>Project web: http://arda.cern.ch. The middleware studies can be found under http://lcg.web.cern.ch/LCG/activities/arda/middleware.html. A list of presentation and written documents is under http://lcg.web.cern.ch/LCG/activities/arda/documents.html</p>				
Milestones for Next Quarter			Status	Comments
<i>ID</i>	<i>Date</i>	<i>Milestone description</i>		
Comments and Additional Information				
<p>Please note that a new set of milestone is being discussed between LCG and the experiments</p>				

QUARTERLY STATUS REPORT				
Project Name			Date	
Distributed Database Deployment			21-Jan-2006	
Report Period			Author Name	
2005 Q4			D. Duellmann	
Milestones for the Quarter			Status	Comments
DBS-1	30.11.05	Hardware setup defined and proposed to the GDB	Done	http://agenda.cern.ch/fullagenda.php?ida=a045327
Summary of Progress				
<p>The project defined after workshop held in October (http://indico.cern.ch/conferenceDisplay.py?confId=a055549) the setup for a first distributed production phase (March-Sept '06) and a proposed workplan for 2006 (approved by GDB and MB). The h/w and s/w setup at the participating sites is reviewed during the weekly project meetings and is progressing well. Several sites have their database clusters already installed and are doing already acceptance tests (BNL, CNAF, IN2P3), other sites (RAL, GridKA) are awaiting the delivery of database hardware. The database, FroNtier and squid installation at tier 0 is also progressing well. The squid server setup for use from FroNtier has been coordinated by CMS at 10 tier 1 and 2 sites. First tests with the CMS software framework (using POOL/FroNtier/squid) in the 3D testbed server show promising results for cached data. Also the Oracle streams based replication tests (online-offline and tier0-tier1) of ATLAS, CMS and LHCb in the 3D testbed are continuing with positive results.</p>				
Outstanding Issues since Last Report				
None				
Milestones Changes and Actions				
References and Hyperlinks				
Project wiki pages with meeting minutes, workplan etc: https://uimon.cern.ch/twiki/bin/view/PSSGroup/LCG3DWiki				
Milestones for Next Quarter			Status	Comments
DBS-2	31.01.06	Hardware setup, acceptance tests and RAC setup - Applies to CERN and all Tier-1 sites participating in the March 06 production: ASGC, BNL, CNAF, FNAL, GridKA, IN2P3 and RAL . The remaining tier 1 sites PIC, NIKHEF/SARA, NDGF and TRIUMF may participate but are only expected to fully join for full production in Sept 2006.	in progress at several sites	status of requested h/w setup available from the 3D wiki.
DBS-3	31.01.06	Experiment software framework includes the tested s/w and data definition for applications which will be deployed during SC4 (in particular the data definition for the subdetector conditions). Applies to ATLAS, CMS and LHCb	in progress	tests with experiment framework and some subdetectors have been performed with COOL by LHCb and with POOL/FroNtier by CMS. ATLAS is planning to move framework based tests.
DBS-4	01.02.06	Tier-1 readiness workshop - readiness reports from Tier-0 and Tier-1 sites about database and squid cache installation. Readiness report from ATLAS, CMS and LHCb about conditions implementation for main subdetectors.	prepared	Upcoming workshop agenda http://agenda.cern.ch/fullagenda.php?ida=a058495

DBS-5	28.02.06	Apps as released before the workshop are installed on the Tier 0 database and connected via streams or FroNtier to the Tier 1 sites (milestone for CERN and early Tier 1 sites). Tier-0 database is populated and replicas / caches are tested by jobs running on tier 1 and tier 2 (milestone for ATLAS, CMS and LHCb)		
DBS-6	31.03.06	Tier-1 services starts - milestone for early production Tier 1 sites		Level 1
Comments and Additional Information				

QUARTERLY STATUS REPORT				
Project Name			Date	
Grid Deployment Board			1-Feb-2006	
Report Period			Author Name	
2005 Q4			Kors Bos	
Milestones for the Quarter			Status	Comments
GDB1	11/10/05	Accounting mini workshop	Done	http://agenda.cern.ch/fullAgenda.php?ida=a055283
GDB2	12/10/05	GDB meeting in Bologna	Done	http://agenda.cern.ch/fullAgenda.php?ida=a055283
GDB3	09/11/05	GDB meeting at CERN	Done	http://agenda.cern.ch/fullAgenda.php?ida=a045327
GDB4	14/11/05	T0/T1 network meeting in Seattle	Done	http://agenda.cern.ch/fullAgenda.php?ida=a056251
Summary of Progress				
No meeting was organised in December on request because of the late date and other commitments.				
Outstanding Issues since Last Report				
See for action items the GDB wiki				
Milestones Changes and Actions				
None.				
References and Hyperlinks				
GDB wiki.				
Milestones for Next Quarter			Status	Comments
GDB5	11/1/05	GDB meeting at CERN		http://agenda.cern.ch/fullAgenda.php?ida=a057701
GDB6	31/1/05	T0/T1 network meeting in Amsterdam		http://agenda.cern.ch/fullAgenda.php?ida=a068
GDB7	8/2/05	GDB meeting at CERN		http://agenda.cern.ch/fullAgenda.php?ida=a057702
GDB8	8/3/05	GDB meeting at CERN		http://agenda.cern.ch/fullAgenda.php?ida=a057703
Comments and Additional Information				

QUARTERLY STATUS REPORT				
Project Name			Date	
ALICE			11-Jan-2006	
Report Period			Author Name	
2005 Q4			Yves Schutz	
Milestones for the Quarter			Status	Comments
MS5	Sep-05	Start of distributed analysis	Delayed	An advanced prototype will be released end of January and tested by selected users.
MS6	Sep-05	Metadata prototype ready	Delayed	The physics metadata have been defined and a prototype of the query system has been tested. Not yet released. Not a critical milestone.
MS7	Dec-05	Condition infrastructure deployed	Done	The framework for calibration and alignment has been released. The implementation by detector groups is in progress.
MS8	Dec-05	Preliminary implementation of algorithms for alignment and calibration ready for all detectors	Delayed	Only 20% of the detectors so far have implemented the calibration and alignment framework. This is a critical milestones which must be ready before the DC6 starts.
Summary of Progress				
<p>Testing of the ALICE Computing Model has continued during the Physics Data Challenge 05 executed together with the LCG Service Challenge 3. The central and distributed Grid services agreed by the Baseline Services group have been adopted. LFC and FTS were partially tested, extensive testing of CASTOR2 and RB has been performed.</p> <p>A complete set of requirements and a simple installation methods of the ALICE VO software on the VO-box has been established and the scalability and robustness of the VO-box model demonstrated. MC generation and reconstruction have been continuously run during 4 months, involving 4 T1s and 10 T2s. A maximum of 2450 concurrent jobs has been reached, evenly distributed over the participating T1s and with a good T2 involvement. The job failure rate was low: below 2% for the application software, below 15% non-persistent failures of Grid services.</p> <p>The majority of produced data has been stored in CASTOR2@CERN, which showed a good stability and continuous operational improvement over the 4 month running period. xrootd has been extensively used as a file transport protocol.</p> <p>The application software (ROOT, AliRoot) was debugged in details and optimized for CPU and memory consumption.</p>				
Outstanding Issues since Last Report				
<p>The frequent updates of the middleware and services prevented more computing centres to be joined to the PDC'05, in addition this introduced unevenness of resource utilization. In terms of operation, priority for updates was given to the centres in the SC3 plan. Not all services, in particular LFC and FTS, could be tested in depth. Not all sites could be switched to the VO-box model of operation, mainly because of lack of experts, who could do the installation of the software (in the absence of full documentation). CASTOR2 is still missing a few functionalities which existed in CASTOR. An uniform SRM functionality across platforms and back-ends is badly needed allowing general SE tests, especially for end-user data-analysis. A higher-level services is needed for FTS. Not all centers can cope with the high demands of the application in terms of memory utilization: ALICE requires worker nodes with at least 2GB memory.</p>				
Milestones Changes and Actions				
<p>The coming milestones address the preparation of the final version of the application software to be used for data processing in year 1, and the preparation of the last data challenge DC06 before data taking to run together with the LCG service challenge 4. Intermediate milestones (not in the TDR) have been set as listed below:</p> <p>MS 100 January 2006: Rerun of SC3 disk – transfers (max 150MB/s). We should get ready to do this with the current data triggered via jobs or scheduled transfers</p> <p>MS 101 February 2006: Start of distributed analysis evaluation by selected users</p> <p>MS 102 March 2006: T0-T1 “loop-back” tests at twice the nominal rate (CERN). We run our bulk production and send data back to CERN</p>				

MS 103 March 2006: PROOF ready at CAF
 MS 104 April 2006: T0-T1 disk-disk (nominal rates) disk-tape (50-75MB/s). First chance to push out data, reconstruction at CERN
 MS 105 April 2006: First tests with PROOF at CAF
 MS 106 July 2006: T0-T1 disk-tape (nominal rates). Second chance to push out the data. Reconstruction at CERN and remote centers.
 MS 107 July 2006: T1-T1, T1-T2, T2-T1 and other rates according to the rated defined in the TDR.
 MS 108 September 2006: Scheduled analysis challenge in T1s.
 MS 109 September 2006: Unscheduled challenge in T2s.

References and Hyperlinks

Milestones for Next Quarter			Status	Comments
MS9	Jan-06	AliRoot release for PDC06 with advanced prototype of alignment and calibration framework, alignment and calibration algorithms for all detectors and global inter-detector alignment	Delayed	A release is still scheduled in January but without complete implementation of detectors calibration algorithms. The release for PDC06 will be done in March before the start of the data challenge.
MS10	Jan-06	Start of PDC06	Delayed	Delayed until March 2006.

Comments and Additional Information

QUARTERLY STATUS REPORT				
Project Name			Date	
ATLAS			30-Jan-2006	
Report Period			Author Name	
2005 Q4			Dario Barberis	
Milestones for the Quarter			Status	Comments
LCG 1.6.10.2	9/30/2005	ATLAS: Use the gLite middleware (version 1.1) on the extended prototype (eventually the pre-production service) and provide feedback (technical issues and collect high-level comments and experience from the experiments)	Done	Testing started only for the few components of gLite that were released to the experiments. Current tests of the new ATLAS Distributed Data Management System include the LFC catalogue and the FTS file transfer service. The WMS is under test by the CERN/ARDA and Milan groups.
LCG 1.1.3	2/15/2005	ATLAS integration: SEAL Integration into Athena	Done	Many SEAL components are already integrated into ATLAS. These include the LCGDict dictionary, and the PyLCGDict2 Python bindings, components of the maths library including Minuit, and the plug-in manager, although integration of the latter is incomplete. Use of the SEAL plug-in manager within Gaudi remains incomplete, although the ATLAS Geant4 system uses it extensively. The completion of this milestone is now included in the plan to adapt to the forthcoming LCG Applications Area reorganization
ATL-EDM-05-01	10/31/2005	Review of the implementation of the Event Data Model for reconstruction.	In progress.	Event Management Board set up to control changes and implement schema evolution.
ATL-DDM-05-01	10/31/2005	Distributed Data Management system ready for operation with LCG Service Challenge 3.	Done	DDM was functional for the Tier-0 tests and the Tier-0 – Tier-1 data transfers.
ATL-CSC-05-01	11/30/2005	Start of simulation production for Computing System Commissioning.	Delayed to April 2006.	The required functionality for CSC tests (complete as-installed detector geometry and full usage of the COOL conditions database for calibrations and alignments) will be available with software release 12 (end of March 2006).
Summary of Progress				
<p>Most attention in Q4 2005 was dedicated to LCG Service Challenge 3. We ran in the SC3 context two major operations: data transfer from Tier-0 to all ATLAS Tier-1s, and internal Tier-0 data flow as specified by the ATLAS computing model.</p> <p>Both operations resulted in the identification (and resolution) of several problems and considerable progress was achieved. Further tests will take place in the context of SC4 in 2006.</p>				
Outstanding Issues since Last Report				
<p>Most open issues are related to the organisation of distributed systems (Data Management, Production and Analysis). The single architectural point that is still open, but to be sorted out in early 2006, is the issue of the deployment of experiment-specific services at Grid sites (VO Boxes).</p>				

Milestones Changes and Actions				
<p>The internal planning of software releases was modified in order to increase the priority of having a complete description of the detector "as built" and "as installed" in the pit. This includes macroscopic geometrical displacements from the nominal positions and breaking the symmetry of the detector and the magnetic field. The planning for ATLAS Computing System Commissioning is now being integrated with the SC4 schedule.</p>				
References and Hyperlinks				
Milestones for Next Quarter			Status	Comments
ATL-REL-06-01	1/31/2006	Production release for Computing System Commissioning and early cosmic-ray studies; completion of the implementation of the Event Data Model for reconstruction.	Delayed to end March 2006	See above comments.
ATL-CSC-06-01	2/28/2006	Start of Computing System Commissioning.	Work in progress.	Now coordinated with SC4 activities. Tier-0 tests and production operations that do not need SC4 infrastructure will take place starting in March 2006.
Comments and Additional Information				

QUARTERLY STATUS REPORT				
Project Name			Date	
CMS			20-Jan-2006	
Report Period			Author Name	
2005 Q4			Lothar A. T. Bauerdick	
Milestones for the Quarter			Status	Comments
CPT-4204	16-Sep-05	Computing systems ready for Service Challenge SC3	done	
CPT-4212	16-Dec-05	Computing systems ready for Cosmic Challenge	delayed	see below
Summary of Progress				
<p>The main milestone for this quarter was to be ready for the Service Challenge 3, and to run SC3 in phases.</p> <p>CMS had specified as its main goal, to establish the "baseline use case" (described in detail in the C-TDR) of moving data from the T0 tape system to T1 centers, onto tape for "custodial storage". After "selection jobs" would run on this data at each T1 center, "selected data" would be moved to T2 centers for fake "data analysis" jobs. At the same time, "simulated data" from T2 centers would move to the T1 centers for storage and further distribution. Also exchange of datasets between T1 centers was planned for "service phase 2", but this was de-scoped due to the problems encountered en suite of SC3.</p> <p>An initial assessment of the CMS results from SC3 are given here: https://uimon.cern.ch/twiki/bin/view/CMS/CmsServiceChallengeInitialAssessment</p> <p>The new Data Management system with its components Dataset Bookkeeping System (DBS) and Dataset Location Services (DLS) was scheduled to be initially deployed in November, but due to manpower shortage (in particular loss of the key developers) this was delayed to early 2006. This eliminates any contingency on delivering the MC production system and data management system in preparation for SC4.</p> <p>The new MC production system that allows to run MC with the re-engineered framework was designed and the prototype development was started, with the initial release now scheduled for early 2006</p> <p>The re-engineered EDM and framework is being delivered: The re-factoring of the simulation, reconstruction, event filter and visualization Software (CMSSW) in the context of the new framework and event data model is well underway. The release of version 0.3.0 will be the base for the Cosmic Challenge data-taking. This release provides a fully functional system, from the readout of the detector up to local reconstruction of nearly all the sub-detectors, and Geant4 simulation.</p>				
Outstanding Issues since Last Report				
<p>This quarter saw a lot of discussion of the role of application specific services at regional centers and the use of VO boxes to operate these services. CMS has written a "White Paper" on the issue, see https://uimon.cern.ch/twiki/bin/view/CMS/CmsServicesAtTier1</p>				
Milestones Changes and Actions				
<p>Comments on the delay of milestone CPT-4212: Given the delay of the start of the CMS Cosmic Challenge running to April 2006, the CPT-4212 milestone was delayed to allow a reasonable assessment of the requirements on the services required from CMS Computing and the bookkeeping interfaces.</p> <p>The underlying services for file transfers from SX5 to the CERN computing center, the required hardware resources and the system to move data to interested institutions outside CERN are existing.</p> <p>The Cosmic Challenge running will proceed during April and the computing systems will be ready by then. In the mean time a CMS Computing Cosmic Challenge Taskforce has been formed to provide the required services in time and to support eventual needs for detectors starting up data taking between now and the start of the challenge.</p>				
References and Hyperlinks				

Milestones for Next Quarter			Status	Comments
CPT-212/C	Mar-06	Computing Systems Ready for Cosmic Challenge	on track	CMS computing services for MTCC are being setup by the C4 Task Force led by T. Wildish
CPT-402/C	Apr-06	Computing Systems Ready for SC4		lead for planning for SC4 is I.Fisk. Currently working out scope and schedule for SC4, agreeing on middleware stack, deployment of CMS application services (DBS, DLS, MCprod) commencing
CPT 404/C-	Mar-06	Computing Systems ready for start of simulation		initial prototype of Mcprod system being prepared, led by P.Elmer
Comments and Additional Information				

QUARTERLY STATUS REPORT				
Project Name			Date	
LHCb			6-Jan-2006	
Report Period			Author Name	
2005 Q4			Nick Brook	
Milestones for the Quarter			Status	Comments
	Nov-05	Example of sub-detector alignment	Completed	Draft LHCb notes exists for the VELO alignment. Strategy presented at dedicated alignment workshop in Glasgow & at the December LHCb collaboration meeting
	Nov-05	Analysis at all Tier-1's	Completed	Stripped DSTs distributed to Tier1 centres. Analysis jobs run at all Tier-1's via DIRAC using LCG RB matching sites. Limitations due to lack of disk SE with SRM at some sites
	Dec-05	End of review of recons s/w	Completed	Implementation already underway
Summary of Progress				
Major DIRAC review Implementation of reviewed event model underway (using libCore and new Linear Algebra) Conditions Framework functional and released				
Outstanding Issues since Last Report				
Lack of functionality in Grid m/w - limitation achieving metric in SC3 Delay to s/w releases through AA math library delays				
Milestones Changes and Actions				
Like to repeat part of SC3 to see if one can achieve metrics				
References and Hyperlinks				
https://uimon.cern.ch/twiki/bin/view/LHCb/LHCbTaskForce https://uimon.cern.ch/twiki/bin/view/LHCb/DiracReview https://uimon.cern.ch/twiki/bin/view/LHCb/EventModelReview https://uimon.cern.ch/twiki/bin/view/LHCb/LHCbPhysicsEventModelTaskForce				
Milestones for Next Quarter			Status	Comments
	Mar-06	Final alignment Strategy		
Comments and Additional Information				

