



LHC Computing Grid Project

Quarterly Status and Progress Reports

2006 Q1

11 May 2006

Table of Contents

WLCG High Level Milestones.....	3
Service Challenges.....	5
Grid Sites	
1. ASGC	7
2. CC-IN2P3	9
3. CERN	11
4. FZK.....	15
5. INFN	19
6. NDGF	23
7. PIC	25
8. RAL.....	29
9. SARA-NIKHEF	33
10. TRIUMF	35
11. US ATLAS	37
12. US CMS.....	39
Areas and Projects	
1. Applications Area.....	41
2. Deployment Area	45
3. ARDA.....	49
4. Distributed Databases	51
5. Grid Deployment Board.....	53
Experiments	
1. ALICE	55
2. ATLAS	57
3. CMS.....	59
4. LHCb.....	61

QUARTERLY STATUS REPORT					
Project Name			Date		
WLCG					
Report Period			Author Name		
2006 Q1			Les Robertson		
Milestones for the Quarter					
ID	Date	Milestone	Coord.	Status	Comments
SC4-1	28.02.06	All required software for baseline services deployed and operational at all Tier-1s and at least 20 Tier-2 sites	J.Shiers	Completed	Baseline services installed and operational at all Tier-1s and more than 20 Tier-2s was achieved by the end of SC-3. These have been upgraded with the LCG 2.7 distribution at the end of January.
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	Completed	FNAL, SARA, IN2P3, TRIUMF, BNL and CNAF. Only two sites, CNAF and FZK, on GEANT as anticipated in the last quarterly report.
SC4-2	28.02.06	Use cases and service level support defined for SC4 SC workshop at CHEP	J.Shiers	Completed	Defined by the outcome of the Mumbai workshop, the agreed Tier-0/Tier-1 bandwidth targets, and the site reliability measurement process (agreed at MB on 7 March).
CAS-1	15.03.06	Castor2 Readiness Review	A.Cass	Postponed to early June	The review was re-scheduled for April, but the chair of the review resigned. A new chair has been appointed and new dates agreed.
SC3-4	31.03.06	All services on all Tier-1 sites monitored	J.Shiers	Postponed to end April	Set of tests agreed, but milestone postponed to end April (MB of 7 March). See also GD-5 and GD-6 in the Deployment Milestones plan
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	Completed	Agreed on algorithm for calculating availability and problem response times at MB of 4 April. No changes in the levels set in the MoU Annex 3. Algorithm confirmed at MB of 4 April (Rome).
Summary of Progress					
<ul style="list-style-type: none"> • The Mumbai workshop, attended by about 150 people including representation of experiments, sites and software providers, achieved its goals of agreeing on the services to be provided for SC4, the timetable, and the experiment use cases. • The process for coordinating the testing of the SRM 2.1 implementations for the three mass storage systems used in LCG (Castor, dCache, DPM) was agreed, with a target for production by the end of September. • The service developed for SC3 continues in operation, with an average of 24K jobs per day accounted to LHC throughout the quarter (13.8K per day at CERN+Tier-1 sites, 10.3K at other sites in the EGEE grid). • The Castor 2 system has performed well during intensive Tier-0 testing at CERN during the quarter, both in terms of performance and reliability, with the data recording milestones (CERN-DR-2, CERN-DR-5) easily reaching their targets. An ATLAS Tier-0 test demonstrated their nominal throughput, driven by a 320 MB/s input stream, though only sustained for a 24 hour period. 					
<ul style="list-style-type: none"> • The final part of SC3, the re-run of the CERN-Tier-1 data distribution tests, took place in January, achieving the target of sustained 1 GB/s. • Interoperability between OSG and EGEE is progressing, and some of the resources in OSG are now visible on EGEE and vice versa. • About half of the functionality of SEAL has been completely migrated to ROOT, and the experiments and Applications Area projects have made considerable efforts to integrate these new packages. • The Applications Area software is being adapted for the AMD64 architecture. 					

Outstanding Issues since Last Report					
<ul style="list-style-type: none"> • It was agreed at the Mumbai workshop that the introduction of SRM 2.1 would be delayed until after SC4, due to availability of the dCache implementation, the need for a common agreement on how to define storage classes, and to allow a thorough test plan including interoperability testing to be implemented. A coordination team has been established and this should propose a plan of work with milestones. • Progress has been slower than planned in certifying the SC4 middleware distribution, gLite 3.0. The SC4 planning foresees a 6-week period for testing by experiments on the Pre-Production Service (starting 15 March) followed by one month for installation and set-up at the participating sites. By the end of the quarter the package was still not available on the Pre-Production Service for testing, and re-scheduling will be necessary if the start date of the SC4 service phase (milestone SC4-4: 1 June) is to be maintained. 					
<ul style="list-style-type: none"> • The formal constitution of the Nordic Data Grid Facility (NDGF) has not been completed, and so there is no official management contact between LCG and NDGF. • The full 2006 capacity goals were originally scheduled for the beginning of April. It was agreed in the MB that this date should be moved to 1 September. Progress looks good for building up processor and tape system capacity at most sites, but the installation of disk capacity is going much more slowly. A review of experiment needs for disk storage at Tier-1 centres during the second half of the year is needed. 					
Milestones Changes and Actions					
<ul style="list-style-type: none"> • There have been discussions with each of the experiments on their architectures and plans for DAQ-Tier-0 integration and testing. The current status was presented to the MB on 4 April. While high-level architectures and plans are available, there are many details to be sorted out, and the internal planning of the experiments is not compatible with the targets of LCG milestones CERN-DTT-1 and CERN DTT-2. New milestones will have to be defined to enable us to follow progress in this fundamental area. 					
References and Hyperlinks					
Milestones for Next Quarter					
ID	Date	Milestone	Coord.	Status	Comments
SC4-3	30.04.06	Service Challenge 4 Set-up: Set-up complete and basic service demonstrated, capable of running experiment-supplied packaged test jobs, data distribution tested.	J.Shiers		This should follow a period when the services are available on the Pre-Production Service for testing by the experiments. Delays in integration and certification of the gLite 3.0 release make this date uncertain.
DRC-3	30.04.06	1.0 GB/s data recording demonstration at CERN: Data generator → disk → tape sustaining 1.0 GB/s for one week using the CASTOR 2 mass storage system and the new tape equipment.	B.Panzer	Completed	See CERN-DR-5. 1 GB/s data recording rate sustained for one week, with 1.6 GB/s sustained over a 24 hour period. Completed on March 18.
SC4-4	31.05.06	Service Challenge 4: Start of stable service phase,	J.Shiers		Including all Tier-1s and 30 Tier-2 sites The service must be able to support the full computing model (use cases and services to provide) of each experiment, including simulation and end-user batch analysis at Tier-2 sites. This is dependent on SC4-3, which is uncertain. See "Outstanding Issues".
Comments and Additional Information					

QUARTERLY STATUS REPORT				
Project Name			Date	
Service Challenge 4				
Report Period			Author Name	
2006 Q1			Jamie Shiers	
Milestones for the Quarter			Status	Comments
SC3-NW-1	31.01.06	Demonstrate stable disk - disk transfers between CERN and all Tier1 sites at either the nominal data rate for that site, or 150MB/s (whichever is the lower)	Completed.	CERN configuration was limited to 10Gbps, shared with other activities. All sites managed to achieve their targets, with some exceeding them.
SC3-NW--2	28.02.06	Demonstrate disk - tape transfers at 50MB/s between CERN and 6 Tier1 sites.	Completed.	
SC4-NW-1	01.04.06	Demonstrate disk - disk transfers from CERN to all Tier1 sites at the full nominal rate, average over 24 hour periods. Any drop in data rate by 10% must be both accounted for and made up.	Completed.	80% of the full nominal data rate out of CERN was achieved during the second week, with 6 sites reaching their full nominal rate. Investigation of the rates to other sites is on-going.
SC4-NW-2	01.04.06	Demonstrate disk - tape transfers to all sites at the nominal rate, scaled to the performance of current hardware (50-75MB/s)	Completed.	One site (FZK) unable to participate due to schedule. Transfer rates in general at or close to targets.
Summary of Progress				
<p>The Service Challenge 3 throughput re-runs – both disk-disk and disk-tape – were successful in their objectives of demonstrating improved rate and stability. In individual tests, many of the sites demonstrated rates at or even above their nominal rates. Whilst the steps required to ramp-up these data rates to the full nominal values at all sites – for disk-tape transfers – adding the additional complexity of full T0 and T1 activities in parallel – should not be underestimated, this work nevertheless represents an important milestone in the preparations for full-scale LHC Computing services.</p>				
Outstanding Issues since Last Report				
Milestones Changes and Actions				
References and Hyperlinks				
Milestones for Next Quarter			Status	Comments
SC4-NW-3	31.05.06	All T1 sites to define channels to all other T1s and supported T2s and demonstrate functionality of transfers between sites.	On-going	

SC4-NW-4	31.05.06	Understand issues related to data transfers to those sites for which the full nominal rates were not achieved in April.		
SC4-NW-5	30.06.06	Perform disk - disk transfers to all Tier1 sites of one experiment (ATLAS) under realistic operation conditions (file sizes, transfer rates to individual sites, concurrency both for CERN export buffers and Tier1 site input buffers).		
SC4-NW-6	31.07.06	Perform disk - tape transfers to 6 Tier1 centres at full nominal rates for the sum of the VOs supported by the site in question.		
Comments and Additional Information				

QUARTERLY STATUS REPORT								
Project Name					Date			
ASGC					10.05.2006			
Report Period					Author Name			
2006 Q1					Jason Shih			
Milestones for the Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
	01.01.06			100				
ASGC-11	30.01.06	Complete T2 SC testing and start peak transfer tests					complete	details of network throughput testing have been updated on SC wiki page (APROC)
ASGC-12	30.01.06	Complete 16TB disk procurement					complete	add two disk server to share loading from prod. Castor server.
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)					complete	SRM, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES
ASGC-13	03.02.06	Start operation of expanded STK, disk pool of CASTOR MSS (ver. 1, 36TB)					complete	initialized all disk cache with xfs filesystem type to have better performance.
ASGC-14	15.02.06	16 TB Disk Operational					complete	new disk cache have been separated into 3 physical volumes. One of them have been attached to disk server, and will have rest of two online soon (due to the delay of ongoing procurement of disk servers, we're confirming final spec with local vendors)
ASGC-15	15.02.06	298	39	70	24		complete	network capacity at: 2Gbps
ASGC-16	15.05.06	4 new tape drives (LTO3) and 700 LTO3 cartridges implemented					delayed	as noted in summary, due to the delay of internal procedures. We have rescheduled the delivery of new tape system till mid of May.
ASGC-18	22.05.06	new Tape system, 280T, in operation					delayed	plan to start in operation one week after delivery
MOU-1	01.04.06	950	400	100	500	75	moved to 01.07.06	from WLCG MoU do not edit
ASGC-19	30.04.06	CASTOR2 Data migration done (phase 1)					delayed	details are put in summary
ASGC-20	20.03.06	add new CE dedicated for biomed DC2					complete	subcluster connecting to new CE have provide another 90 KSI2k dedicated for biomed DC2
ASGC-21	01.04.07	start operation of new CE					complete	validating with SFT and start taking biomed jobs at same time.
ASGC-22	20.04.06	MDM demonstration					in progress	complete testbed installation, update from MDM will help installing needed s/w here.
ASGC-23	10.05.06	3D testbed					complete	ASGC help setting up RAC testbed, and have start streaming with CERN, as well as completed initial staged for this project.
Summary of Progress								
<p>Mid March, VOBOX installation have been completed, and atlas DDM experts help troubleshooting with DQ2 (prod ver) installed at ASGC. Planning to upgrade to 0.2.7 in near future, when receive confirmation from Atlas DDM group.</p> <p>Due to the delay of tape expansion procurement internally, we've postpone new tape system implementation to mid of May while openbid will of current procurement will be completed first week of May.</p>								

Removed action item #10, postpone new core router installation, as well as loop migrate. Besides, item 17 have been deleted as well.

For the castor2 installation and data migration, we have late in obtaining LSF license from platform (CASTOR2 could only accept LSF as batch for data management requests. Rescheduled deliverable of castor2 platform ready end of Apr. and start in operation with new tape system added, which will be mid of May.

Ongoing biomed DC2 activities (with wisdom and DIANE framework) mainly adopt computing resource from new introduced subcluster, as well as 90 processors from production farm.

May 11, new tape system plan to installed this Fri., and will connect to SC castor stager servers for functionality testing. If all testings goes well, will continue with disk to tape throughput testing with CERN.

VO Boxes Installations at the End of the Quarter								
VO	Status	Comments						
ALICE								
ATLAS	installed	vobox have been installed since mid of March, and deployment of D2 for Atlas DDM is also available on it, including basic throughput testing comparing between glite IO and FTS.						
CMS								
LHCb								
Capacity Available at the End of the Quarter								
Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments		
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)			
	480	36	100	23	10			
Outstanding Issues since Last Report								
Milestones Changes and Actions								
References and Hyperlinks								
MDM Proj link: http://lists.grid.sinica.edu.tw/taskforcewiki/MDM_Demo_Project_2006								
ASGC 3D proj link: http://lists.grid.sinica.edu.tw/apwiki/Build_Oracle_RAC_10g_Release2_Cluster_on_IA64_and_Scientific_Linux_CERN_43								
Milestones for Next Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
SC4-3	30.04.06	SC4: Set-up complete and basic service demonstrated					complete	SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES
ASGC-24	14.05.06	SRM/FTS throughput validations with T2s					in progress	real data mgmt testing with T2 centers in AP
ASGC-25	24.05.06	LCG production farm upgrade to glite 3.0					in progress	confirm progress in WLCG meeting, will complete the upgrade in two weeks
ASGC-26	20.05.06	split CE from batch server					in progress	this stabilize the scheduling functionality and also help improving availability of ASGC. IS have strong
ASGC-27	15.06.06	replacing old tower boxes with 1U server (for LCG server nodes)					in progress	confirming spec with vendor, prefer having 20 ibm 306m 1u servers with dual core cpu model, besides, IMPI
ASGC-28	01.07.06	Start procurement of CPU, Disk and Tape					in progress	confirming spec with vendor, should turn to old H model of IBM blades in this expansion.
MOU-1	01.07.06	950	400	100	500	75		from WLCG MoU do not edit
Comments and Additional Information								

QUARTERLY STATUS REPORT									
Project Name					Date				
CC-IN2P3					11/04/2006				
Report Period					Author Name				
2006 Q1					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-05-6	31.12.05	Evaluation process for compute nodes and disk servers completed					Complete	The evaluation process finished by the end of January 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					Ongoing	Received 2 IBM tape drives for evaluation	
IN-06-3	17.01.06	Dedicated network link Lyon-CERN (10 Gbps) operational					Complete		
IN-06-4	31.01.06	End of the evaluation process for compute nodes					Complete	The result of this evaluation was used for specifying the technical needs for the call for tender for compute nodes.	
IN-06-5	15.02.06	Begin formal procurement process for compute nodes					Ongoing	The EU wide call for tender is running. Officially published on March 4th. Deadline: April 24th.	
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES	
IN-06-6	28.02.06	Additional tape drives installed and operational					Complete	10 additional STK 9940 tape drives added to the HPSS production setup.	
IN-06-7	28.02.06	End of the evaluation process for disk storage hardware					Complete	The result of this evaluation was used for specifying the technical needs for the call for tender for disk servers.	
IN-06-18	28.02.06	Middleware LCG 2.7.0 deployed on all services.					Complete	CE, SE, BDII, FTS, LFC, ... upgraded to LCG 2.7.0	
IN-06-8	15.03.06	Begin formal procurement process for disk storage					Ongoing	The EU wide call for tender is running. Officially published on week 14th. Deadline: May 22nd.	
MOU-1	01.04.06	1170	520	200	535	75	moved to 01.07.06	from WLCG MoU do not edit	
Summary of Progress									
<p>Procurement: The efforts on this quarter were devoted to prepare and launch the procurement procedures for compute nodes and disk servers.</p> <p>Hardware Upgrades: the MSS production setup was upgraded and the project of the evolution of the automated cartridge library started. The dCache setup was also upgraded for preparing both for SC3 and for SC4.</p> <p>Connectivity: The 10 Gbps dedicated link to CERN is operational since mid-january and was used for throughput tests during SC3.</p> <p>Middleware Deployment: Deployment of LCG 2.7.0 was done in February and the work on the pre-production testbed continued.</p> <p>Service Challenges: during SC3 throughput tests rerun in January, the target rate of 150 MB/sec for CERN(disk)→CC-IN2P3(disk) transfers was sustained during ~48 hours. Target rate of 50 MB/sec for CERN(disk)→CC-IN2P3(tape) exceeded during 4 days.</p> <p>Project Management: one person was appointed to coordinate the technical activities of the LCG-France project related to the Tier-2 and Tier-3 sites, in particular the interactions with the French Tier-1. This work is ongoing.</p> <p>Human Resources: a senior engineer hired as a dedicated support for Atlas at the French Tier-1. She started working on March 1st. Currently: we have 3 people dedicated to the support of the 4 LHC experiments.</p>									

VO Boxes Installations at the End of the Quarter									
VO	Status	Comments							
ALICE	Operational	Delivered to experiment on October 2005 (see QR for 2005Q4).							
ATLAS	Operational	Delivered to experiment on October 2005 (see QR for 2005Q4).							
CMS	Operational	Delivered to experiment on October 2005 (see QR for 2005Q4).							
LHCb	Operational	Delivered to experiment on October 2005 (see QR for 2005Q4).							
Capacity Available at the End of the Quarter									
Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments			
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)				
92.27%	660	87	150	215	50	Note that the unavailability fraction (7.73%) was scheduled and published as service downtime.			
Outstanding Issues since Last Report									
<p>For administrative reasons related to the hardware procurement procedures, the availability of additional compute nodes and disk servers will be delayed. The additional computing capacity, initially planned to be operational by the end of May is rescheduled to mid-July. The additional disk storage capacity, initially planned to be operational by the end of June is rescheduled to mid-August and mid-October.</p> <p>The works for upgrading the machine room, initially planned to be finished by the end of May are re-scheduled in two phases: in June, a first phase for partially upgrading the cooling capacity and a second phase in September for a) finalising the upgrade of the cooling capacity and b) performing the upgrade of the power capacity. The second phase is expected to last until the end 2006 and will involve some scheduled service interruptions.</p>									
Milestones Changes and Actions									
The milestones related to the capacity increase (CPU and disk) are re-scheduled (see comments above).									
References and Hyperlinks									
Milestones for Next Quarter									
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments	
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)			
MOU-1	01.04.06	1170	520	200	535	75	moved to 01.07.06	from WLCG MoU do not edit	
SC4-1	30.04.06	SC4: Set-up complete and basic service demonstrated							SRM, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES
IN-06-14	15.04.06	LCG-3D: migration of LHC data bases (FTS & LFC) to the local Oracle cluster.							The local FTS and LFC grid services for LHC experiments use the Oracle cluster as the backend. Some of those databases will be used to test replication between Oracle clusters of Tier-1s.
IN-06-16	31.05.06	Upgrade of production grid services (LFC, FTS, CE, SE, ...) to the appropriate version of gLite performed.							All baseline services ready for SC4.
IN-06-9	31.05.06	Upgrade of machine room power and cooling finished					Rescheduled		The work for upgrading the machine room infrastructure will start on June and last until end of 2006.
IN-06-10	31.05.06	800			400		Rescheduled	The increase of compute capacity for LHC experiments will be done by mid-July 2006, instead of end May.	
IN-06-13	15.06.06	LCG-3D: interconnection of Oracle cluster with CERN's cluster for replication of LHC data bases. Initial tests performed.							
IN-06-12	15.06.06	Tests of data transfer between Tier-1 and 3 Tier-2s performed using FTS and SRM end-points.							
IN-06-17	15.06.06	Modify configuration of HPSS-backed dCache for sustaining nominal rate to tape of 200 MB/sec during throughput tests in July 2006.							
IN-06-15	30.06.06	LCG-3D: initial tests of monitoring of interconnected Oracle databases with CERN.							
IN-06-18	30.06.06	Improve monitoring of the site's grid services and progressively integrate them into the existing operational procedures for 24x7 service of the whole site.							
IN-06-11	30.06.06		300				Rescheduled	The additional storage capacity will be installed during August and October 2006.	
Comments and Additional Information									
I would propose to replace the title of the column 'Tape' by 'MSS' because this term includes also the disk used as cache of data physically on tape.									

QUARTERLY STATUS REPORT								
Project Name					Date			
FZK / GridKa					9-May-2006			
Report Period					Author Name			
2006 Q1					Holger Marten			
Milestones for the Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
	01.01.06			150				
GK-06-1	31.01.06	Network: Start of testing lightpath to CERN					shifted to 31.03.06	Waiting for DFN/GEANT2 progress
GK-06-2	31.01.06	Start procurement for expansion of tape library and tape IO with 4 drives LTO-3					done	
GK-06-3	31.01.06	Start procurement for expansion of Write Pool with 100 MB/s (file servers)					done	
GK-06-4	31.01.06				395	40	done	Tape: Expansion with 250 TB of existing 145 TB Tape IO: Expansion with 4 drives LTO-3 of existing 8 drives LTO- 2
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)					done	SRM, LFC, FTS, CE, RB, BDII, RGMA, VO BOXES
GK-06-5	Feb/Mar 06	Network: Lightpath to CERN; testing, configuration; performance & stability checks					removed	
GK-06-6	Feb/Mar 06	Tape: Tape access & I/O optimization tests (GridKa internal)					done	
GK-06-7	Mar 06	CPU: delivery, installation, configuration, tests					in pogress	CPU: Expansion with 500 kSI of existing 530 kSI Procurement process running; precise date of delivery to be defined after selection of contractor
GK-06-1	31.01.06 new date Mar 06	Network: Start of testing lightpath to CERN					shifted to April 06	Waiting for DFN/GEANT2 progress
GK-06-8	21.03.06	Pre-production services installed and ready for testing by experiments					in progress	
GK-06-9	31.03.06	Network: Additional 10 Gbps OPN GridKa-CERN operational					shifted to May 06	Network: 2x 10 Gbps for 2006 complete and operational

OPN-2	31.03.06	Tier-0/1 high performance network operational					delayed	
MOU-1	01.04.06	1030	280	200	393	200	moved to 01.07.06	from WLCG MoU do not edit

Summary of Progress

FZK/GridKa successfully contributed to the SC3 re-run in January 2006, with peak data rates disk-to-disk in access of 200 MB/s.

FZK/GridKa also participated in the Mumbai SC4 workshop, with several contributions at the CHEP 2006 and in the VO-Box Meetings in January and April 2006.

A combined Tier1-Tier2 meeting was held in combination with the 47th session of the GridKa Technical Advisory Board on March 2-3 in Karlsruhe.

A second machine hall with sufficient electrical power and cooling capacity has been refurbished for the installation in 2006. Racks for 2006 have been installed.

Remarks on milestones

GK-06-3 File servers to expand the dCache WritePool have already been delivered and are currently installed for SC4 throughput tests.

GK-06-4 This milestone completes a procurement of LTO-3 drives and tape media started at the end of October 2005, and the successive installation thereof. This milestone completes the tape pledged for LCG at GridKa for 2006.

SC4-1 All listed services have been installed except for VO-Boxes of CMS and LHCb. Detailed setup of these are still under discussion.

GK-06-6 This milestone solves an issue found during SC3 concerning the connection between dCache and the underlying TSM for tape access. Namely, after each file transfer to tape, the dCache/TSM interface sent a signal "transfer finished" which initiated TSM to dismount the tape even if the same medium could in principle be used for the next transfer, thus largely degrading the overall throughput to tape. The interface has been rewritten and tested.

GK-06-7 (CPUs at GridKa)

150 new machines (dual core, dual Opteron) have been delivered and mounted into the racks in calendar week 12 (end of March). Software installation, configuration and testing is scheduled for April. Consequently, the 1030 kSI2k envisioned for 3.4.2006 will only be achieved one month later. However, due to the procurement procedures, no better planning was available before.

GK-06-8 (preproduction installation)

Due to the delay in delivery of gLite 3.0 the pre-production installation conflicted with a larger network maintenance at GridKa on March 20-22 and with an already scheduled mw installation support of the German ATLAS Tier-2 in Munich, and could thus not be finished in time. In progress, rescheduled for April 10.

VO Boxes Installations at the End of the Quarter

VO	Status	Comments
ALICE	installed	LCG 2.7
ATLAS	installed	LCG 2.7 + gLite functionalities
CMS	in progress	PhEDEx up and running; DB proxy under discussion
LHCb	in progress	

Capacity Available at the End of the Quarter

Availability (%)	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Comments
	530	83	200	395	75	

Outstanding Issues since Last Report									
GridKa OPN delayed (see "Milestone Changes and Actions").									
Milestones Changes and Actions									
<p>GK-06-1; GK-06-5; GK-06-9 (OPN at GridKa) The original idea of this set of milestones was as follows: Delivery of the OPN connection by DFN/Geant in January (GK-06-1), a configuration and testing phase of two months in February/March (GK-06-5) and readiness of this final Tier0-Tier1 connection by end of March (GK-06-9) for usage during the SC4 throughput tests in April.</p> <p>Although already shifted to March 06, the additional OPN connection has not yet been delivered, such that GK-06-1 and GK-06-9 had to be shifted by an additional month to April and May, respectively. It has thus been decided to continue running the throughput tests in April via the already existing 10 Gbps Internet connection as a fall back, and to switch to the OPN connection as soon as it becomes available. The original, additional "safety factor" of two months for testing the OPN (GK-06-5) has been removed.</p> <p>GK-06-14 (disk at GridKa) This milestone was shifted by one month after first discussions with the selected provider for the disk in 2006. However, this does not influence the high level milestone to have all upgrades available at GridKa on first of July.</p>									
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	
GK-06-10	03.04.06	1030	83	100	395	40	Available 04/2006	CPU: 1030 kSI2k for 2006 complete and operational Tape: 393 TB for 2006 complete and operational	
GK-06-8	21.03.06 new date 10.04.06	Pre-production services installed and ready for testing by experiments							
GK-06-7	Mar 06 new date Apr 06	CPU: installation, configuration, tests							CPU: Expansion with 500 kSI of existing 530 kSI CPUs have been delivered and mounted in the racks.
GK-06-1	31.01.06 new date Mar 06 new date Apr 06	Network: Start of testing lightpath to CERN							Waiting for DFN/GEANT2 progress
GK-06-9	31.03.2006 new date May 06	Network: Additional 10 Gbps OPN GridKa-CERN operational						Available 05/2006	Network: 2x 10 Gbps for 2006 complete and operational

GK-06-11	30.04.06	Write Pool (server): Expansion with 100 MB/s to 250 MB/s						Tape IO: Expansion of existing 12 drives with 4 drives LTO-3
GK-06-12	30.04.06			200		200		
SC4-3	30.04.06	SC4: Set-up complete and basic service demonstrated						SRM, LFC, FTS, CE, RB, BDII, RGMA, VO BOXES
GK-06-13	May-06	Installation of gLite 3.0 in production environment						
SC4-4	31.05.06	SC4: Start of stable service phase						
GK-06-14	May-06 new date Jun-06	disk: delivery, installation, configuration, tests						Disk: Expansion with 200 TB of existing 80 TB Procurement process running; precise date of delivery to be defined with contractor
GK-06-15	May 06	start procurement of tape robotics for 2007+						
GK-06-16	Jun 06	start procurement of racks for hardware 2007						
MOU-1	01.07.06	1030	280	200	393	200	from WLCG MoU do not edit	
GK-06-17	01.07.06	1030	280	200	395	200	Available 07/2006 Disk: 280 TB for 2006 complete and operational	
GK-06-18	31.07.06	Analysis & classification of Tier-1 services concerning 24x7 availability and system stability available for GridKa						
GK-06-19	Sep 06	start procurement of LAN hardware and cabling for 2007						
GK-06-20	30.09.06	Concept for for 24x7 operations (on-call services, shifts) available for GridKa						
DBS-8	30.09.06	Full Database service in place - milestone for all Tier-1 sites						Level 1
GK-06-21	30.09.06	Full Database services in place at GridKa						
GK-06-22	Oct 06	delivery, installation and test of tape robot for 2007+						details to be defined after selection of provider
GK-06-23	Oct 06	Discussions with FZK administration concerning 24x7 operations cost, contracts etc.						
GK-06-24	Nov-06	Concept for for 24x7 operations discussion in GridKa OverView Board						
GK-06-25	Nov-06	Delivery of racks for hardware 2007; installation for chilled water and electricity						
GK-06-26	Dec 06	Installation of LAN cabling for hardware 2007						
Comments and Additional Information								

QUARTERLY STATUS REPORT								
Project Name					Date			
INFN					10.05.2006			
Report Period					Author Name			
2006 Q1					Tiziana Ferrari, Luca dell'Agnello			
Milestones for the Quarter								
ID	Date	CPU	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					done	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
	31.12.05	1. Upgrade to 10 Gb/s of the entire network path connecting CERN and CNAF completed 2. Testing of the new path completed					done	schedule shifted to the 1st week of March due to hardware fault problems, that have been fixed by replacement of the NICs on the local core switch, and small routing problems, that have been fixed
	31.12.05	Addition of new 160 TB of disk space, to be put in production according to the requirements of the experiments					done	
	16.01.06			150				
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES
CN-06-1	01.02.06	Possibile upgrade to CASTOR v2					done	testing phase started in the 1st half of Dec. Deployment on going. We foresee to use CASTOR v.2 for the throughout phase of SC4 (April 06)
CN-06-2	01.02.06	New tapes purchased: 120 TB					on going	schedule shifted to the end of March due to delays with the delivery of the hardware
CN-06-3	01.02.06	1500 KSI2k	235 TB	150 MB/s	320 TB		done	2500 CPU slots (out of 1000 physical CPUs) LHC: up to 2500 CPU slots (Worker Nodes are shared)
CN-06-4	01.02.06	Revision of local and national connectivity					done done done	* Usage of policy routing at the GARR access point * Type of connectivity to INFN Tier-2 under discussion • Backup link Tier-1 <-> Tier-1 (Karlsruhe) under discussion (see CN-06-13)
CN-06-5	15.02.06	Upgrade of CERN-CNAF dedicated link to 10 Gbps					done	done during 1st week of March
CN-06-6	01.03.06	Evaluation of dCache and StoRM (for disk-only SRMs)					done	we still need to choose the technical solution to be adopted in production. Tests on StoRM will continue during the next two Quarters, as new features and development work are expected by the end of June.

CN-06-7	31.03.06	start of test phase of LCG3d with LHCb and ATLAS					done	the start time of these tests is scheduled on the second half of April
SC-04-1	15.03.06	Middleware and services available on the pre-production service for beta testing by the experiments SC4: installation of the services from gLite 3.0 in PPS at CNAF completed					done	most of the PPS gLite 3.0 services requested at CNAF have been installed, with the only exception of FTS and of the RGMA server (see new Milestone for details)
MOU-1	01.07.06	1800	850	200	850	200	moved to 01.07.06	from WLCG MoU do not edit

Summary of Progress

- **Procurement:** The efforts on this quarter were devoted to prepare and launch the procurement procedures for compute nodes and disk space.
- **Hardware Upgrades:** the MSS production setup was upgraded (addition of three tape drives) The Castor setup was also upgraded in preparation to SC4. Upgrade to Castor2 to version 2.0.3 for production services is scheduled at the end of April.
- **Connectivity:** The 10 Gbps dedicated link to CERN, based on the GEANT2 L1/L2 service, is operational since the first week of March and is going to be used for the pre-SC4 disk-disk throughput tests. Extensive testing of the connection has been carried out in Jan; the network path exhibits high reliability and good overall performance.
- **Middleware Deployment:** Deployment of LCG 2.7.0 was done at the end of March, and the installation work on the pre-production testbed is almost completed.

VO Boxes Installations at the End of the Quarter

VO	Status	Comments
ALICE	installed	Installation available since SC3
ATLAS	installed	Installation available since SC3
CMS	installed	Installation available since SC3
LHCb	not available	Installation requested during the 2nd half of March. This request has been left pending in order to wait for general guidelines from the task force devoted to this. In the meanwhile, the decision of INFN is to avoid the installation of VO Boxes.

Capacity Available at the End of the Quarter

Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)	
	1600	235	84	320	not tested yet	Performance to tape is going to be tested in April (pre-SC4 testing)

Outstanding Issues since Last Report

- Castor2:** Castor2 upgrade to 2.0.3 is needed to improve the Castor2 performance problems experienced during the pre-SC4 throughput tests (April 2006). This is needed in order to put the Castor2 service in production.
- FTS:** the installation of FTS with mysql backup in the PPS was not possible in March, due to outstanding bugs. The installation of FTS with an Oracle backend was hence decided. Because of manpower issues, this installation has been postponed to the first half of May (see new Milestone).
- Pre-SC4 Throughput tests:** given the poor pre-SC4 disk-disk throughput performance recorded by CNAF because of local Castor2 problems, we foresee the possibility to re-run the throughput tests during the first half of May.

Milestones Changes and Actions								
New milestone (CN-06-11) added for Q2 about the deployment of Castor2 in production.								
New milestone (CN-06-12) added about the installation of a new FTS instance with Oracle backend in the PPS.								
Network connectivity: new milestones added about the implementation of the backup connection to FZK and about the T1-T2 connectivity at INFN.								
References and Hyperlinks								
Service Challenge pages at INFN: http://grid.infn.it/scinfn								
Milestones for Next Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
CN-06-9	10.04.06	purchase of 3 new drives for STK library					done	Already used in production and for the service challenge tests
CN-06-10	30.04.06	Deployment of the LHC 3D DB services and infrastructure: setup of LFC replication service (in the framework of LCG3d project)					done	testing phase starting in May in collaboration with LHCb group.
CN-06-11	30.04.06	Deployment of the LHC 3D DB services and infrastructure: setup of Oracle clusters for LHCb and ATLAS					on going	Clusters have been installed but not yet used. LHCb cluster will host the production read-only replica of LFC
CN-06-12	01.05.06	Procurement of additional 66 TB of disk space					on going	66 TB (amount actually available to users)
CN-06-13	15.05.06	Deployment of Castor2 for production services					on going	
SC-04-1	15.05.06	Middleware and services available on the pre-production service for beta testing by the experiments SC4: installation of the services from gLite 3.0 in PPS at CNAF completed: Installation of FTS with Oracle backend in PPS; installation of R-GMA server					on going	
CN-06-14	15.05.06	Re-run of CERN-CNAF disk-disk throughput tests (target: 200 MB/s)					on going	
CN-06-15	30.06.06	Implementation of CNAF-FZK backup connection via cross-border fibres, involving GARR, DFN and SWITCH					on going	
CN-06-16	30.06.06	Implementation of new CNAF-GARR uplink (10 GigaEthernet) for a bandwidth upgrade of the existing 1 GigaEthernet connection supporting T1-T2 connectivity					on going	
CN-06-17	30.06.06	Testing of new StoRM features					on going	

CN-06-18	30.06.06	Testing of trasfers Tier1 - Tier2: testing of latest gLite FTS server release by ATLAS in collaboration with various ATLAS Tier-2 sites at INFN (Milano, Napoli, Roma)						
CN-06-19	30.06.06	Preparation work for support 24x7 from Jan 2007						
MOU-1	01.07.06	1800	850	200	850	200	from WLCG MoU do not edit	
SC4-3	30.04.06	SC4: Set-up complete and basic service demonstrated						SRM, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES
Comments and Additional Information								

QUARTERLY STATUS REPORT								
Project Name					Date			
NDGF								
Report Period					Author Name			
2006 Q1					Michael Grønager			
Milestones for the Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
	01.01.06	410	88	150	130	75		
	28.02.06	Investigation of dCache for DPM					Completed	For SC3 we were using DPM and dCACHE seems more usable
	01.03.06	New AS number to make NDGF appear as one network for CERN					Postponed	All sites will use an IP in the new AS range
	01.01.06	361	26	150	85	-		No tape test so far. See comments...
NDGF-SC3	16.01.06	NSC, PDC and NBI will participate in the SC3 rerun using the successful DPM setup from summer 2005					Completed	
LCGARC-3	18.01.06	ARC-LCG interoperability meeting in Uppsala					Completed	
T1COORD-2	18.01.06	Coordination meeting in Uppsala.					Completed	
HWNBI-4	31.03.06	NBI compute and storage system in production					Partly	LoadLeveler Interface not yet done. Postponed to 30.05.06
NETWORK-1	15.02.06	Network coordination meeting					Completed	
NDGF-SC4-1	15.02.06	Initiate dCache test setup					Completed	
LCGARC-3	23.03.06	gLite RB patch for LCG->ARC demonstrated					Partly	
LCGARC-4	24.03.06	ARC-LCG interoperability task-force meeting at CERN					Completed	
DCACHE-1	27.03.06	dCache Single DB, distributed Pools installed					Completed	A more well suited setup was used
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					Postponed	Postponed to Q3 2006
	31.03.06	505	52	150	225	75		
MOU-1	01.04.06	520	160	150	240	75	moved to 01.07.06	from WLCG MoU do not edit
Summary of Progress								
<p>SC3: Participated successfully in the SC3-rerun with a distributed setup similar to the setup from SC3. 4 Sites took part in the SC3.</p> <p>SC4: Design of the distributed storage setup completed in collaboration with DESY: Each site run its own dCache version. For SC4 a central dCache streams the data to the site dCaches for LHC the central service will be replaced with an SRM-proxy which is under development</p> <p>NETWORK: Internal Nordic network has been designed and procured</p> <p>HARDWARE: Storage and CPU funding approved in 3 countries (awaits 1) for 2007</p> <p>INTEROP: Interoperability between ARC and LCG is on track and is coordinated with EGEE-SA3. Other interoperability services for ALICE and ATLAS has been identified and are under testing.</p>								

VO Boxes Installations at the End of the Quarter									
VO	Status	Comments							
ALICE									
ATLAS									
CMS									
LHCb									
Capacity Available at the End of the Quarter									
Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments			
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)				
	505	52	150	225	75	Copenhagen site still needs a loadleveler interface to the			
Outstanding Issues since Last Report									
<p>The creation of the LoadLeveler-Grid interface to the Copenhagen site has been delayed to May 2006. Interoperability between ARC and gLite has been delayed – awaits availability of more stable gLite3.0. The move to a Nordic OPN has been postponed to Q3 2006 as more sites needs integration.</p>									
Milestones Changes and Actions									
References and Hyperlinks									
See the NDGF wiki for further details : www.ndgf.org/wiki (login can be obtained from waananen@nbi.dk)									
Milestones for Next Quarter									
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments	
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)			
MOU-1	01.04.06	520	160	150	240	75	moved to 01.07.06	from WLCG MoU do not edit	
NDGF-SC4-2	03.04.06	SC4 T0-T1 disk-disk transfer test begins							
NDGF-SC4-3	24.04.06	SC4 T0-T1 disk-tape transfer test begins							
VOBOX-2	15.05.06	ALICE-ARC VOBOX installed also at NBI							
GLITE-1	01.05.06	Upgrade of services to gLite3 begins							
DCACHE-2	30.05.06	Distributed dCache setup ready							
LCGARC-3	22.05.06	gLite RB patch for LCG->ARC demonstrated							
LCGARC-5	02.06.06	ARC-LCG interoperability task-force meeting at Helsinki							
HWNBI-4	30.05.06	NBI compute and storage system in production							LoadLeveler Interface not yet
NETWORK-2	Q3	Upgrade to use a single address space for Nordic sites, upgrade to use dedicated connections from sites to geant – completed							
	01.05.06	Target for full participation in SC4							
Comments and Additional Information									
<p>Currently intensive planning is taking place, with meetings regarding:</p> <ul style="list-style-type: none"> * LCG-ARC interoperability setup * The ARC Alice VO Box * ARC gateway to the ATLAS DDM * General setup of the Tier-1 storage services (including 24x7 plan) <p>However, a more thorough plan is not to be expected before earliest in 2 weeks. So any further immediate additions to the QR would be on too loose ground.</p>									

QUARTERLY STATUS REPORT								
Project Name					Date			
PIC								
Report Period					Author Name			
2006 Q1					Gonzalo Merino			
Milestones for the Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
PIC-05-7	01.12.05	FTSserver-atlas installed					pending	Not yet deployed. Experiment need delayed.
PIC-05-12	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.
2006								
	01.01.06			100		na		
PIC-06-1	01.01.06	start deployment of FTS server with MySQL backend					done	
PIC-06-2	15.01.06	start deployment of Castor2 test service					done	
PIC-06-3	15.01.06	New IBM tape library delivered					done	
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)					ongoing	SRM, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES
PIC-06-4	01.02.06	FTS server deployed					postponed to 21/04/06	Problems with FTS-MySQL. Had to switch to FTS-Oracle.
PIC-06-5	01.02.06	start testing d-Cache for deploying an SRM-disk service					delayed to 28/02/06	Late incorporation of new engineer dedicated to this project.
PIC-06-6	01.02.06	start tender for disk expansion					postponed to 1/04/06	Vendor delay on SAS availability.
PIC-06-7	01.02.06	start deployment of additional 1Gbps WAN infrastructure					delayed to 3/04/06	Delay by Catalan regional academic network.
PIC-06-8	01.03.06	new 2Gbps WAN infrastructure operational					delayed to 15/05/06	Due to PIC-06-7
PIC-06-9	01.03.06	150	41.5	60	85	60	delayed to 15/05/06	Due to PIC-06-7
PIC-06-10	01.03.06	place order for disk expansion					delayed to 1/05/06	Due to PIC-06-6
PIC-06-11	01.03.06	1st version of SRM-dCache available for testing					delayed to 15/06/06	Partly due to PIC-06-5, partly due to problems encountered and poor documentation.
PIC-06-12	07.03.06	Castor2/SRM2.1 service ready for testing					delayed to 15/06/06	Error in estimating the effort needed. Additional complexity of adopting two new proprietary technologies: LSF and Oracle.

PIC-06-13	15.03.06	new IBM tape library integrated within Castor					postponed to 15/05/06	No manpower could be allocated for this task since it was all taken by Castor2 deployment.
PIC-06-14	01.04.06	SRM-dCache deployed in production					delayed to 1/07/06	Due to PIC-06-11
PIC-06-15	01.04.06	start disk expansion deployment					delayed to 1/06/06	Due to PIC-06-6
PIC-06-16	01.04.06	Start of SC4-setup phase						
PIC-06-17	01.04.06	Add tape cartridge capacity					done	
PIC-06-18	01.04.06	150	41.5	60	90	60	delayed to 15/05/06	The tape capacity (90TB) has been delivered, but the WAN has been delayed due to PIC-06-7
MOU-1	01.04.06	250	140	100	158	75	moved to 01.07.06	from WLCG MoU do not edit

Summary of Progress

One of the major activities planned for this quarter was the deployment of a new Castor2 instance. Originally, the driver for this was that SRM2.1 was needed for SC4 startup on 1st April. After the Mumbai workshop, the SRM2.1 constraint moved towards October 2006. The Castor2 deployment at PIC started around mid January and is still ongoing. The main reason for this delay has been that the effort needed was underestimated, specially since it made us to adopt two proprietary technologies (namely Oracle and LSF) which were completely new at PIC. An activity also launched during last quarter was the test of dCache as a possible option for providing: a) an SRM-disk service and b) a disk pool management system able to deal with insecure disks (such as secondary disks from WNs). The work had to start almost one month later than initially foreseen, due to the late incorporation to the team of the engineer allocated for this task. Once the work could start, we could see that we had underestimated the learning curve for dCache, so we have to delay this milestone one extra month.

The FTS server deployment started at the beginning of January. We chose to deploy the MySQL version first because at that time no Oracle know-how existed at PIC and second because we wanted to try whether if the performance was enough. Soon some serious functionality problems were spotted, and finally after interacting with the developers, their recommendation was to switch to Oracle. The FTS-Oracle deployment started then on mid March and is currently under certification by the local ATLAS team.

The disk capacity procurement has been delayed two months, mainly due to the fact that we are interested in purchasing SAS technology, and the provider is delayed in releasing this new product. On the other hand, we have not seen the need of purchasing disk capacity with the old technology to meet the cMoU milestone for 2006 on 1st April, since the requirements from the experiments at the time of writing this report are not reflecting the need for such capacity in the short term (only ATLAS is requesting 3TB of disk for the 2nd quarter 2006, according to <https://twiki.cern.ch/twiki/bin/view/LCG/SitePIC>).

Regarding SC4-1 milestone, on the target date all of the services were deployed but the FTS and the CMS and LH

VO Boxes Installations at the End of the Quarter

VO	Status	Comments
ALICE	n.a.	PIC does not support ALICE
ATLAS	done	Deployed since november 2005. Reported in Q42005.
CMS	ongoing	Currently working with the CMS-PIC contact person to move Phedex box into vobox-cms. Milestone PIC-06-26 for Q2-2006.
LHCb	ongoing	Currently working with the LHCb-PIC contact person to deploy the vobox-lhcb. Milestone PIC-06-25 for Q2-2006.

Capacity Available at the End of the Quarter

Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)	
	150	41.5		90		

Outstanding Issues since Last Report								
<p>The deployment of an extra dedicated 1Gbps WAN link, expected to start in february, has a delay of more than one month due to the late delivery from the regional academic network provider. A consequence of this is that we have started the SC4-throughput phase still with the old 1Gbps shared WAN infrastructure.</p> <p>The current status of the deployment of dCache and Castor2 at PIC still does not allow us to have a clear idea on which technology we shall use to implement the SRM-disk service requested by experiments. Our current plan is to continue with the deployment and testing of these two systems during the next quarter and reach some conclusion with the experience gained. For the moment, and until these activities are advanced enough at PIC, we are planning to provide both the SRM-tape and SRM-disk using Castor1.</p> <p>Another important open issue for us is to understand which technology might allow us to use the local disk space in the farm nodes in a consistent way. We are currently testing the dCache Resilient Manager to evaluate up to which point it can solve this problem.</p> <p>The extra manpower needed for the Castor2 deployment has had an impact in delaying the deployment of a new IBM</p>								
Milestones Changes and Actions								
<p>Following the request from LHCB, the vobox-lhcb is planned to be deployed by the 3rd week of April. For CMS, we are running a Phedex-box since long time and we foresee we will move this service inside a vobox-cms by 1st May 2006.</p> <p>The deployment of an FTS server has been postponed until the 21st April.</p> <p>The deployment of the additional 1Gbps dedicated network link has been postponed until the 21st April.</p> <p>The deployment of an SRM-disk service based on dCache and a Castor2 SRM-tape service for testing have been postponed until the 1st May.</p> <p>The integration of the new IBM tape library within Castor has been postponed until mid-May.</p> <p>The disk expansion milestones have been postponed two months, so we plan to deploy the disk capacity needed to reach the 77 TB on July the 1st.</p>								
References and Hyperlinks								
Milestones for Next Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
PIC-06-14	01.04.06	SRM-dCache deployed in production					delayed to 1/07/06	Due to PIC-06-11
PIC-06-15	01.04.06	start disk expansion deployment					delayed to 1/06/06	Due to PIC-06-6
PIC-06-16	01.04.06	Start of SC4-setup phase						
PIC-06-17	01.04.06	Add tape cartridge capacity					done	
PIC-06-18	01.04.06	150	41.5	60	90	60	delayed to 15/05/06	The tape capacity (90TB) has been delivered, but the
MOU-1	01.04.06	250	140	100	158	75	moved to 01.07.06	from WLCG MoU do not edit
SC4-3	30.04.06	SC4: Set-up complete and basic service						
PIC-06-25	21.04.06	vobox-lhcb deployed						
PIC-06-10	01.05.06	place order for disk expansion						Delayed from previous Quarter
PIC-06-4	15.05.06	FTS server deployed						Postponed from previous Quarter
PIC-06-13	15.05.06	new IBM tape library integrated within Castor						Postponed from previous Quarter
PIC-06-26	15.05.06	CMS Phedex box moved into vobox-cms						

PIC-06-8	15.05.06	new 2Gbps WAN infrastructure operational						Delayed from previous Quarter
PIC-06-9	15.05.06	150	41.5	60	90	60		Delayed from previous Quarter (same as PIC-06-18)
PIC-06-12	15.06.06	Castor2 service ready for testing						Delayed from previous Quarter
PIC-06-31		T1-T2 transfer test milestone for ATLAS-T2s						On 08/05/2006: waiting for definition from experiment contact persons
PIC-06-32		T1-T2 transfer test milestone for CMS-T2s						On 08/05/2006: waiting for definition from experiment contact persons
PIC-06-33		T1-T2 transfer test milestone for LHCb-T2s						On 08/05/2006: waiting for definition from experiment contact persons
PIC-06-15	01.06.06	start disk expansion deployment						Delayed from previous Quarter
PIC-06-14	01.07.06	SRM-dCache deployed in production						Delayed from previous Quarter
PIC-06-19	01.07.06	disk expansion deployed						Delayed from previous Quarter
PIC-06-20	01.07.06	150	77	60	90	60		Delayed from previous Quarter
MOU-1	01.07.06	250	140	100	158	75		from WLCG MoU do not edit
Comments and Additional Information								
<p>We just want to state again that the capacity milestones for 2006 were requested to be modified in the cMoU for PIC into: cpu/disk/tape = 140/77/89. We would like to see these numbers in these sheets instead of the current 250/136/158.</p>								

QUARTERLY STATUS REPORT								
Project Name					Date			
RAL								
Report Period					Author Name			
2006 Q1					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-16	01.01.06			150		na		
RAL-5	15.11.05	Commence Service Hardening Project					Started	
RAL-10	15.12.05	Tier-1 10Gbit uplink switch ordered					Completed	
RAL-11	15.12.05	Delivery of 3D Production Hardware					Completed	Delivered Feb 2006. Order was delayed because of contention
2006								
RAL-16	01.01.06			150		na		
RAL-17	15.01.06	CASTOR2 Hardware Delivered					Completed	
RAL-18	15.01.06	Delivery of Site Edge Router 10Gbit Upgrade					Completed	
RAL-19	15.01.06	4*6TB disk resource allocated to dteam in preparation for SC3 throughput test					Completed	resource was allocated in March. Not all capacity has yet been released by the experiments
RAL-20	15.01.06	On-Call System in Place					Started	
RAL-21	15.01.06	Airconditioning Capacity Upgrade Installed					Completed	Completed in February
RAL-22	15.01.06	dCache Upgraded to version 1.6.6 (SRM 1)					Completed	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					Not Started	Depends on timely FTS release. No release available yet.
RAL-24	15.01.06	dCache ready for 150MB/s disk to disk.					Completed	
RAL-25	15.01.06	Tier-1 10Gbit uplink switch delivered					Completed	
RAL-26	15.01.06	493	120	150	229	75		
RAL-27	31.01.06	50% of Tier-2 GRIDPP sites tested at 300-500Mb/s for 1TB bi-directional from tier-1					Started	About 40% of sites tested.
RAL-28	31.01.06	Migration complete from old to new Robot					Completed	Functionality only
RAL-29	31.01.06	Test CASTOR2 System Running					Completed	Depends on effort from CERN
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)					Completed	SRM, 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)					Not Started	Depends on receiving dCache SRM 2.1 on 15th January and deploying 2-3 weeks later. SRM 2.1 support is not yet available in dCache
RAL-31	15.02.06	3D Service Moves to Production Hardware					Started	Scheduled for production in April. Delayed by late delivery of hardware.
RAL-32	15.02.06	Tier-1 Connected to Lightpath Edge Router at 10Gb/s					Completed	
RAL-33	15.02.06	1st Disk and CPU delivery					Completed	Delivered 10th March

RAL-34	28.02.06	New Tape drives Installed					Completed	Not funded by Tier-1 but available for test. T10K drives are operational in CASTOR and also from the old ADS tape service.
RAL-35	28.02.06	UKLIGHT Provisioned at 4*1 Gbit					Completed	May be earlier but scheduled January date now impacted by throughput test. Delivered in February after problems with the NMS.
RAL-36	28.02.06	Order Tape Drives and Media for new Robot					Not Started	Waiting finalised GRIDPP spending plans. Sufficient drives and capacity already exist (from other funding) to carry out service challenge testing. This purchase is intended to regularise the position by taking ownership of the hardware.
RAL-37	15.03.06	Nagios monitoring System deployed (replaces SURE)					Started	A NAGIOS infrastructure is under test but deployment has been delayed by staff illness and other commitments.
RAL-38	31.03.06	Completion of Phase I Service Hardening					Not Started	
RAL-39	31.03.06	Place order for second delivery of disk/CPU					Not Started	Currently finalising spending plans.
RAL-40	31.03.06	100% of Tier-2 GRIDPP sites tested at 300-500Mb/s for 1TB bi-directional from Tier-1					Not Started	We expect to complete 90% of sites by September.
RAL-41	31.03.06	Local (Preliminary) Throughput Test of CASTOR2 Complete.					Completed	Limited throughput test - 50MB/s?
RAL-42	31.03.06	ATLAS, ALICE and CMS VOs up to date.					Completed	
MOU-1	01.04.06	980	450	150	664	75	moved to 01.07.06	from WLCG MoU do not edit
Summary of Progress								
<p>A number of key developments have progressed this quarter. A new 6000 slot SL8500 robot has been installed and successfully moved into production, replacing the old STK silo. New T10K tape drives have been installed and integrated into CASTOR. The CASTOR deployment remains on track. New hardware has been installed for core CASTOR services, a minimal service installed (with the help of the CERN CASTOR team) and by the end of the quarter the full CASTOR chain had been tested right through to the new tape drives. Procurements of capacity disk and CPU had been completed and the hardware delivered, deployment will be completed next quarter. Deployment of the 3D hardware has been completed and the 3D service will go live next quarter. The lightpath to CERN was upgraded to 4*1Gbit, and the Tier-1s connection to the lightpath router was upgraded to 10Gbit. Throughput testing of the Tier-2 sites is progressing with about 40% of sites able to exceed 300Mb/s to/from the Tier-1. The Tier-1 has been able to exceed 1Gb/s aggregate throughput to 4 or more Tier-2s over a period of 2 days - this has been a valuable exercise with a number of problems being identified. CPU efficiencies (cpu time/wall time) are now gathered automatically and are becoming a very useful diagnostic tool.</p>								
VO Boxes Installations at the End of the Quarter								
VO	Status		Comments					
ALICE	In Production							
ATLAS	In Production							
CMS			CMS didn't ask for a VO-box; instead FroNtier Squid for CMS has been					
LHCb	Installed		Handed over to LHCb to deploy their software;					

Capacity Available at the End of the Quarter									
Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments			
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)				
	225	64	150	70	75	This data now reflects actual allocations to LCG - which depends on capacity available and LHC experiment requests. Disk capacity includes dteam allocation but not tape buffers. CPU capacity reflects Q1 allocations to LCG on the MAUI scheduler. Actual demand (176KSI2K) for CPU was somewhat less than available capacity.			
Outstanding Issues since Last Report									
<p>Tier-2 testing has identified an 800Mb/s bottleneck at the RAL site firewall (Netscreen). When outbound rates exceed this rate, the firewall occasionally throttles throughput. At present the RAL site is running with an 800Mb/s rate limiter (note this does not impact the lightpath traffic). Tier-2 testing also highlighted problems with FTS 1.3, during aggregate throughput testing (inbound) instabilities were experienced with FTS which occasionally hung. Problems were also encountered with ORACLE logs rapidly filling up an 80GB disk. An upgrade to FTS 1.4 is scheduled and more tests will then be carried out to see if the problems remain. Although 7 Tier-2 sites have been successfully tested, the remaining sites have struggled to achieve target rates, this is Service hardning is going slower than planned as the hardware installations have taken priority. The lack of FTS support for SRM copy is a serious concern, it will be challanging to meet SC4 target rates without this functionality.</p>									
Milestones Changes and Actions									
<p>RAL-44 is posponed indefinitely as the site security audit has been cancelled. It will be rescheduled once new dates are available. RAL-46 is now expected early in Q3. RAL-48: Total RAL: capacity is expected to climb to 1079KSI2K by May 2006, LCG allocation for Q2 is currently expected to be 279KSI2K in line with actual demand. A further delivery of 500KSI2K and 300TB of disk is expected in time for September.</p>									
References and Hyperlinks									
<p>CPU efficiency data can be found at: http://www.gridpp.rl.ac.uk/stats/#Eff_stats .The GRIDPP Tier-2 testing schedule and results can be found at: http://wiki.gridpp.ac.uk/wiki/Service_Challenge_Transfer_Tests</p>									
Milestones for Next Quarter									
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments	
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)			
MOU-1	01.04.06	980	450	150	664	75	moved to 01.07.06	from WLCG MoU do not edit	
RAL-43	15.04.06	Delivery of (6?) Tape Storage Bricks and Media for Tier-1							It is likely that only 1-3 storage bricks will be purchased in FY06. Both bandwidth and capacity will only be purchased if proven demand exists.
SC4-3	30.04.06	SC4: Set-up complete and basic service demonstrated							SRM, LFC, FTS, CE, RB, BDII, RGMA
RAL-44	30.04.06	Tier-1 ready for Site Security Audit							Posponed indefinitely
RAL-44A	30.04.06	FTS Upgraded to version 1.4.							
RAL-44B	30.04.06	Separate Test FTS deployed for GRIDPP service challenge stress test.							
RAL-45	15.05.06	dteam disk allocation (4*6TB) moved to CASTOR2							

RAL-46	15.05.06	2nd Delivery of Disk, CPU						Orders have been delayed. Delivery now likely in early July. Expect to buy about 500KSI2K of CPU and 300TB of disk.
RAL-47	15.05.06	1st CPU Upgrade In Production						266KSI2K
RAL-48	15.05.06	717	120	150	229	75	WAN 4x1 GB/s.	
RAL-49	15.06.06	Production Service Commences on New Tape Drives						
RAL-49A	15.06.06	10Gb core switch installed for Tier-1 network. Most downlinks now at 10Gb.						
RAL-50	15.06.06	Pre-Production CASTOR 2 ready for Throughput test						
RAL-51	15.05.06	717	120	150	664	75	WAN 4x1 GB/s.	
RAL-52	31.05.06	1st Disk Upgrade In Production						Disk deployment is expected to be earlier than originally planned (130TB)
RAL-52A	30.06.06	Glite 3 release in production						
Comments and Additional Information								
Milestones are not given for software products with no scheduled release date for production deployment. RAL expects to be able to deploy LCG middleware within 2-4 weeks of software becoming available.								

QUARTERLY STATUS REPORT								
Project Name					Date			
SARA-NIKHEF					10.05.2006			
Report Period					Author Name			
2006 Q1					Jeff Templon			
Milestones for the Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
SN-05-4	31.12.05	purchase of 5 general server machines					postponed	see SN-06-08
2006								
	01.01.06			150		na		
SN-06-01	31.01.06	300	10	150	50	45		5 data servers with 2TB each
SN-06-02	06.01.06	Upgrade dCache to 166-3					done	
SN-06-03	15.01.06	dedicated 10 GE link to CERN					done	
SN-06-04	28.02.06	setup T2 tests					postponed	
SN-06-05	28.02.06	separate T1 tape storage from general tape storage					ongoing	new deadline: 1.4.06
SN-06-06	28.02.06	purchase 400 9940B cartridges					ongoing	ordered
SN-06-07	28.02.06	purchase 4 STK9940B tape drives					ongoing	ordered
SN-06-08	28.02.06	purchase general purpose servers					ongoing	ordered
SN-06-09	31.03.06	new SARA 10G internal network equipment					ongoing	ordered
SN-06-10	31.03.06	300	10	150	130	90		now with 6 T1 dedicated tape drives
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM, LFC, FTS, CE, RB, BDII, RGMA,VOBOXES
SN-06-11	31.03.06	new SARA 10G internal network infrastructure					postponed	all traffic through new Cisco switch
SN-06-12	31.03.06	300	10	150	130	90		now with 6 T1 dedicated tape drives
MOU-1	01.04.06	306	170	150	143	75	moved to 01.07.06	from WLCG MoU do not edit
Summary of Progress								
Successful participation in SC4. NIKHEF switched from Quattor config of LCG to YAIM in order to speed upgrades.								
VO Boxes Installations at the End of the Quarter								
VO	Status	Comments						
ALICE	working	at both NIKHEF and SARA						
ATLAS	working	at SARA, our understanding is that one box suffices						
CMS	N/A							
LHCb								

Capacity Available at the End of the Quarter								
Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments		
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)			
	1422	15	130	130				
Outstanding Issues since Last Report								
Frequent SC4 network problems (source now identified and under investigation by Cisco. Unresolved issue with many-to-one SGM mappings on VO boxes.								
Milestones Changes and Actions								
References and Hyperlinks								
Milestones for Next Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
MOU-1	01.04.06	306	170	150	143	75	moved to 01.07.06	from WLCG MoU do not edit
SC4-3	30.04.06	SC4: Set-up complete and basic service demonstrated						SRM 2.1, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES
SN-06-13	31.04.06	purchase 10 new storage servers with 5 TB disk each						
SN-06-14	31.04.06	purchase 400 kSI2k in new worker nodes						replacement of old nodes (100 kSI2k)
SN-06-14a	10.05.06	3D database server installation in progress						
SN-06-15	30.06.06	tape drives and storage servers installed						10 storage servers 4 back-end servers SAN infrastructure 4 tape drives
SN-06-16	30.06.06	600	84	450	75	90		We were asked to comment here on the disk discrepancy here and MoU. Our best estimate is that the experiments total request is 63 TB in October.
SN-06-17	31.12.06	24 x 7 progress milestone						This milestone was requested by LCG to show progress of 24 x 7
Comments and Additional Information								
3D completion milestone is not possible to plan given uncertainties in Oracle licensing.								

QUARTERLY STATUS REPORT										
Project Name					Date					
TRIUMF										
Report Period					Author Name					
2006 Q1					Reda Tafirout					
Milestones for the Quarter							Status	Comments		
TF-05-5	30.11.05	DWDM optics for 10G TRIUMF/BCNET				done	expected delivery mi-February			
TF-05-7	15.11.05	SC3 phase 2: ATLAS tests, T0-T1 exercise				done	happened only in December of 2005			
TF-05-9	05.12.05	Hardware evaluation/testing: dual-core blades (Dell)				done	3 months evaluation period			
2006										
	01.01.06			50		na				
TF-06-1	15.01.06	SC3 re-run				complete	achieved 140 MB/s sustained			
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)					SRM, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES			
TF-06-2	15.02.06	delivery of DWDM optics for 10 Gig link & optical tests				done	DWDM optics received 3rd week of february			
TF-06-3	01.03.06	Ordering of two tape drives + cartridges				done				
TF-06-4	08.03.06	LCG 2.7.0 site upgrade				done				
TF-06-5	13.03.06	FTS upgrade to version 1.4				done				
TF-06-6	13.03.06	dCache upgrade to 1.6.6				done				
MOU-1	01.04.06	170	27	50	100	50	moved to 01.07.06	from WLCG MoU do not edit		
Summary of Progress										
<p>The main achievements for this quarter were the successful SC3 re-run where we achieved a stable throughput of about 75 MB/s. We were also able to sustain about 140 MB/s when the other sites were not involved and without using FTS. As of this date all necessary electronics, optics and fiber have been purchased to establish a dedicated 10G connection from TRIUMF to Amsterdam. The remaining portion of the link from Amsterdam to CERN is still unavailable. This is expected to be resolved by the early summer. Presently all baseline services for SC4 (ATLAS) are available: LFC, FTS with all channel definitions (Tier-1's and Tier-2's), SRM and VO box for ATLAS. All these services are being maintained and operational.</p>										
VO Boxes Installations at the End of the Quarter										
VO	Status	Comments								
ALICE	NA	experiment not supported								
ATLAS	completed	in service since 2005 / Q4								
CMS	NA	experiment not supported								
LHCb	NA	experiment not supported								
Capacity Available at the End of the Quarter										
Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments				
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)					
98	18	6	75	8.5	30	Presently a modest LCG cluster				

Outstanding Issues since Last Report								
Milestones Changes and Actions								
References and Hyperlinks								
Information on current setup can be found at: http://grid.triumf.ca/status/sc3.html								
Milestones for Next Quarter							Status	Comments
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
TF-06-7	03.04.06	SC4 disk-disk (start)						
TF-06-8	18.04.06	SC4 disk-tape (start)						
TF-06-9	12.05.06	gLite 3.0 deployment						
TF-06-10	15.05.06	10G link tests						between TRIUMF and Carleton University (Ottawa)
TF-06-12	01.06.06	ATLAS SC4 (start)						
TF-06-13	15.06.06	Tier-1 /Tier-2's transfer tests						
TF-06-14	26.06.06	LAN PHY Network module tests						Between TRIUMF and Carleton University in cooperation with Nucleon
TF-06-15	01.08.06	10 Gig link TRIUMF-CERN available						
TF-06-16	15.09.06	DB/ D3D service deployment						
Comments and Additional Information								
Tier-1 got funded only recently, so realistically new hardware will only be available and commissioned toward the end of Q3 of 2006. Milestones with respect to hardware purchasing and computing centre infrastructure installations will be set in the next quarter report. This applies as well for setting milestones for 24x7 support.								

QUARTERLY STATUS REPORT								
Project Name						Date		
BNL - US ATLAS						10.4.06		
Report Period						Author Name		
2006 Q1						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-7	15.12.05	Expanded LAN infrastructure order placed					Complete	Expected to complete by 15-Jan-06
UA-9	31.12.05	OSG 0.4.0 deployed and operational					Complete	Expect to complete by 31-Jan-06
OPN-1	31.12.05	Tier-0/1 high-performance network operational at CERN and 3 Tier 1's						BNL will not be 1 of these 3
2006								
	01.01.06	500	150	150	150	60		
UA-10	15.01.06	Begin installation of expanded LAN infrastructure					Complete	Expect delay to 15-Feb-06
UA-11	15.01.06	Begin installation of new Tape subsystem					Complete	Expect delay to 15-Feb-06
UA-12	15.02.06	Expanded LAN infrastructure operational					Complete	Expect delay to 1-Mar-06
UA-13	15.02.06	New Tape subsystem operational					Delayed	Expected to complete by 14-Apr-06
UA-14	15.02.06	500	150	200	300	200	Delayed	Expected to complete by 14-Apr-06
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)						SRM, LFC, FTS, CE, BDII, RGMA, VOBOXES
UA-15	28.02.06	BNL Wide Area Network Upgrade (OC48 -> 2xλ) operational					Complete	
UA-16	31.03.06	10Gbs BNL Tier 1 <-> Tier 0 dedicated connectivity operational					Complete	
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					Complete	
MOU-1	01.04.06	1120	520	200	300	200	moved to 01.07.06	from WLCG plan do not edit
UA-18	01.04.06	Begin CPU/dCache disk expansion procurement					Delayed	Expect delay to 1-May-06
UA-19	01.04.06	Begin Central disk expansion procurement					Delayed	Expect delay to 1-May-06
Summary of Progress								
<p>During the reporting period all aspects of the BNL Tier 1 network upgrade were completed including both LAN and WAN through to CERN giving the Tier 1 access to total WAN connectivity up to 20 Gb/sec. Also during the reporting period sufficient additional power and cooling capacity was installed to support the current year's planned upgrades of both the US ATLAS Tier 1 and the collocated RHIC computing center.</p> <p>The procurement and installation of a new Tier 1 mass storage system, including servers, disk cache, robotics, tape drives and network connectivity, was also completed. Also during the reporting period a mechanism for automatically delivering CPU accounting information for the US ATLAS Tier 1 to the Grid Operations Centre in the UK was established.</p>								

VO Boxes Installations at the End of the Quarter								
VO	Status	Comments						
ALICE	No							
ATLAS	Operational	Primary use: ATLAS DQ2 Distributed Data Management System						
CMS	No							
LHCb	No							
Capacity Available at the End of the Quarter								
Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments		
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)			
	500	150	200	150	60			
Outstanding Issues since Last Report								
<p>There were two issues during the reporting period. The first was delay in the new tape system and while at the end of the period it was fully installed and under test, it is not expected to be fully operational until the second week of April. This should still be in time for use during the initial tape throughput phase of SC4. The second issue is a delay by one month in the receipt of the funding which will be used for procurements in support of CPU/dCache disk and central disk expansion resulting in delays of associated milestones.</p>								
Milestones Changes and Actions								
<p>There are no changes in the plans. The delay in funding will result in some slippage in follow on milestones as indicated. However, as a result of slack built into the schedule, it is not at this time clear that there will be a delay in the delivery of the fully operational CPU/dCache disk and central disk expansions.</p>								
References and Hyperlinks								
Milestones for Next Quarter								
ID	Date	CPU	Disk	N=>Disk	Tape	N=>Tape	Status	Comments
		(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)		
UA-20	15.04.06	CPU/dCache disk expansion order placed					Delayed	Expect delay to 15-May-06
UA-21	15.04.06	Central disk expansion order placed					Delayed	Expect delay to 15-May-06
SC4-3	30.04.06	SC4: Set-up complete and basic service demonstrated						SRM, LFC, FTS, CE, BDII, RGMA, VOBOXES
UA-22	15.05.06	Begin CPU/dCache disk expansion installation					Delayed	Expect delay to 15-Jun-06
UA-23	15.05.06	Begin Central disk expansion installation					Delayed	Expect delay to 15-Jun-06
SC4-4	31.05.06	SC4: Start of stable service phase						
UA-24	30.06.06	CPU/dCache disk expansion operational						
UA-25	30.06.06	Central disk expansion operational						
UA-26	30.06.06	1125	525	200	300	200		
MOU-1	01.07.06	1120	520	200	300	200		from WLCG plan do not edit
Comments and Additional Information								

QUARTERLY STATUS REPORT								
Project Name					Date			
FNAL - US CMS					01/04/2006			
Report Period					Author Name			
2006 Q1					Ian Fisk			
Milestones for the Quarter								
ID	Date	CPU (kSI2K)	Disk (TB)	N=>Disk (MB/sec)	Tape (TB)	N=>Tape (MB/sec)	Status	Comments
UC-05-3	31.12.05	Hierarchical Fair Share in Condor					Technology Available	(See comment below)
2006								
UC-06-1	01.01.06	1000	100	150	250	na		
MOU-1	01.01.06	728	100	150	250	na	moved to 01.07.06	from WLCG MoU do not edit
UC-06-02	15.01.06	LCG 2.7 Upgrade (CE, SE, BDII,)					Done	(See comment below)
UC-06-03	15.01.06	Upgrade to SRM 2					Initial Deployment (01/04/2006)	(See comment Below)
UC-06-04	01.02.06	RGMA (If Available for Condor)					Not available	(See Comment Below)
SC4-1	31.01.06	SC4: All required software for baseline services deployed (for 28.02.06)					Done	
UC-06-05	01.02.06	SC4 CMS Local Catalog Functionality (LFC, GRLS, Trivial Catalog, etc)					Done	(See comment below)
UC-06-06		Third 6509 Network Switch Acquisition					Ordered Estimated 1/5/06	One site waiting for installation
UC-06-07		30Gb/s Between worker nodes and disk servers					Done	Deployed and well utilized
UC-06-08	01.04.06	1000	96+47 resilient	200	420	7 drives for CMS + 8 shared		
	01.04.06	1000	100	200	300	75		
MOU-1	01.01.06	728	100	150	250	75	moved to 01.07.06	from WLCG MoU do not edit
Summary of Progress								
<p>The US-CMS Tier-1 center has entered the second year of a three year procurement and deployment ramp in preparation for data taking by the start of 2008. We would like to reach 50% complexity for a deployed center by the end of 2006 and we are roughly on schedule. The purchase orders for the first very large procurement of disks have been prepared. The technical evaluations for the second large order of worker nodes has been completed. There is a plan for increased robotic storage involving a new tape room and new robots, which will be executed this summer. This new shared facility will support CDF, D0, and CMS.</p> <p>The grid interfaces both for the OSG and LCG were upgraded during the first quarter. We continue to see high access rate with CMS simulated event production and user analysis jobs. We have begun the process of installing the gLite PPS service. We are also installing services to support CMS including prototypes of the simulated event production infrastructure and new data management components.</p>								
VO Boxes Installations at the End of the Quarter								
VO	Status	Comments						
ALICE								
ATLAS								
CMS	Complete	PhEEx is operational. V2.3 will be installed next quarter						
LHCb								
Capacity Available at the End of the Quarter								
Availability	CPU	Disk	N=>Disk	Tape	N=>Tape	Comments		
(%)	(kSI2K)	(TB)	(MB/sec)	(TB)	(MB/sec)			
	1000	100	200	300	75			

Outstanding Issues since Last Report									
The US-CMS Tier-1 center continues to see a high rate of grid job failures. We have made several architectural changes intended to boost the reliability of the center. A number of failures we have observed are caused by external factors and we look forward to the effort proposed by the LCG-MB in failure diagnosis and resolution.									
Milestones Changes and Actions									
References and Hyperlinks									
The Tier-1 presentation at the 2006 US funding agenda review is available at http://www.uscms.org/SoftwareComputing/ProjectManagement/DOE_NSFReviews/2006-02/talks/Fac_Parallel/Bakken-DOE-Feb606.Tier1Facilities.pdf									
Milestones for Next Quarter									
ID	Date	CPU (kSI2K)	Disk (TB)	LAN=>Dis (MB/sec)	Tape (TB)	LAN=>Tap (MB/sec)	Status	Comments	
MOU-1	01.01.06	728	100	150	250	75	moved to 01.07.06	from WLCG MoU do not edit	
	15.04.06	Tender additional 300 TB Disk and 50 Data Servers (100 Total)					First half order in		
	15.05.06	Tender of 300 Dual Core Dual Opteron Nodes					Eval unit arriving		
	15.05.06	Commissioning of Disk Storage							
SC4-3	30.04.06	SC4: Set-up complete and basic service demonstrated							SRM, LFC, FTS, CE, RB, BDII, RGMA, VOBOXES
	01.06.06	2 Week demonstration of 24/7 Support							
	15.06.06	Production Instance of FronTier from LCG3D							
	01.06.06-15.06.06	Participation in CMS SC4 Activities (Rerun of CMS SC3 Goals with SC4 Infrastructure)							
	01.07.06	1800	700+100 resilient	300	500	7 drives for CMS + 8 shared			
	01.07.06	1790	700	300	500	"7 drives for CMS			
	01.07.06	Second 10Gb/s Research lambda to StarLight							
	01.08.06	Calibration Demonstration with Frontier from LCG3D							
Comments and Additional Information									
<p>Comments regarding the milestones above:</p> <p>Hierarchical scheduling technology was deployed and demonstrated during the first quarter of 2006. It will be enabled by default during the second quarter to balance between local, LCG, and OSG submissions. At a coarse level, each submission interface can be mapped to a group.</p> <p>The upgrade to 2.7 was fairly smooth and also involved removing the shared file system requirements from Grid submissions. Up to this point, we had been relying on a commercial distributed file system. The removal of the shared filesystem has improved our operational stability.</p> <p>SRMV2 was deployed for dCache between March and April of 2006. Improved space reservation is available as well as basic namespace interactions. The interface will be tested during the SC4 throughput challenge</p> <p>The basic connection between RGMA and Condor appears still to be missing. The dynamic information providers for the BDII are available, and the coarse accounting can be provided, but the RGMA functionality is missing.</p> <p>The trivial file catalog functionality has been demonstrated and is in the process of deployment with the new CMS data model. The site configuration file for the trivial file catalog is available, the CMS application that uses the trivial catalog has been released, and the draft name space has been deployed on dCache.</p>									

QUARTERLY STATUS REPORT				
Project Name		Date		
Applications Area				
Report Period		Author Name		
2006 Q1		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.	In progress. Rescheduled to 30.06.06	Work is ongoing to provide the web and distribution kits from the XML description files. Other work of higher priority during the quarter led to a rescheduling of this milestone into the second quarter.
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.	Achieved	The new HyperNews pilot service was presented at an Applications Area meeting in February. After this it was transformed into a production-quality service. It is used by ATLAS and CMS,
SPI-3	31.03.06	Provide the interconnection/interoperability between the savannah and HyperNews services.	Achieved	During discussions with the users it became clear that the functionality of the services are sufficiently distinct such that no direct interoperability between the two services beyond the cross-linking via hyperlinks is required.
SPI-4	31.03.06	Generate CMT configuration and distribution kits from the common (XML based) configuration description.	In progress. Rescheduled to 30.06.06	The implementation of the functionality is done, deployment will take place after a discussion with the maintainers of CMT and the main users which was rescheduled for the second quarter of 2006.
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.	Postponed. Rescheduled to 31.03.07	This milestone is postponed. It can only be done once the CINT data structures are replaced by Reflex. This work is currently on going and is expected to be completed by the end of 2006.
ROOT-8	31.03.06	Have the rootcint dictionary code generator interfaced with the Reflex and gccxml options	Achieved	This milestone has been achieved and the code is part of the version 5.10. Currently 3 options are supported by rootcint -A, rootcint -cint -B, rootcint -reflex -C, rootcint -gccxml The option A is the current default option generating the CINT data structures and API. The option B generates calls to the Reflex API and use an extended version of the CINT header files parser. It will become the default option as soon as the CINT/Reflex integration is completed. The option C generates the same code as option B, but requires the gccxml parser.

COOL-1	30.11.05	Conditions Database (COOL) release based on the latest version of RAL including bulk insertion operations and extended tagging functionality.	Achieved	<p>Internal implementation port from RAL to CORAL achieved in COOL 1.2.7 (January 16). API still based on POOL, with CORAL extensions for early adopters. Extended tagging functionalities ("user tags" and "HVS") achieved in COOL 1.3.0 (April 6). Integration with CORAL connection service functionalities also achieved in COOL 1.3.0 (higher priority for the experiments than multi-channel bulk insertion). New API based on CORAL, no residual POOL dependency.</p> <p>Multi-channel bulk insertion operations delayed to COOL 1.4 (Q2 2006). New milestone proposed.</p>
COOL-3	31.03.06	COOL overall performance study and validation of the experiments requirements. This study should identify the areas that will require further work and optimization.	Achieved	<p>Validation of the ATLAS 'first pass' reconstruction use case achieved in October 2005 (one retrieval of full DCS data snapshots every 5 seconds from an Oracle RAC cluster, corresponding to sustained data rates of 20 MB/s and 200k table rows/s). Performance optimizations for single version folders implemented in COOL 1.3.0 (April 6): removal of linear increase in retrieval time of an IOV with the start-of-validity timestamp. Further performance optimization for multi version folders delayed (no timescale) because the relevant person has left the development team and needs to be replaced. Also identified potential problem with large number of tables in the schema. Will be worked on in Q2-Q3 (no firm milestone as no firm experiment request yet).</p>
SIMU-1	15.12.05	Apply the Fluka-Geant4 (Flugg) geometry interface to one of the LHC calorimeter test-beam simulation.	In progress. Rescheduled to 30.06.06	<p>Technical part of the work completed, consisting in producing hits using Fluka, and then analysing them in the exact same way as for hits produced with Geant4, for the ATLAS TileCal 2002 setup. The work is now devoted to more physics issues: cross check with muons; effect and modeling of the Birks' law; sensitivity of physics observables to the chosen production/tracking cuts; impact of neutrons to the shower shapes. Close interactions with the model developers is now needed and have started. Results are expected for the second quarter of 2006.</p>
SIMU-2	15.12.05	Production quality release of the MC generator level production framework.	Deleted	<p>Original milestone was delayed to 30.06.06 and reformulated as SIMU-8</p>
SIMU-4	31.03.06	First results of the ATLAS combined and 2004 test-beams data comparisons.	In progress. Rescheduled to 30.06.06	<p>Lot of activities are undergoing in ATLAS to analyse the 2004 combined test-beam data, and the first results are now expected now for the second quarter of 2006.</p>

SIMU-5	31.03.06	Monte Carlo event generator files database (MCDB) publicly available and able to deal with large files.	Achieved	Introduction of MCDB Grid certificates and management of large files. Successfully completed on time
SIMU-6	31.10.06	First release of a common framework for handling MC truth information to be used by experiment's simulation programs.	Achieved	Example of MC-truth handling using the HepMC structure has been implemented and communicated to the LHCb and CMS experiments. LHCb is currently implementing a new MC-truth handling based on this example. CMS is evaluating the approach. The example will be also proposed as 'extended example' in Geant4.
SIMU-?	30.04.06	Improved regression suite for release validation and testing infrastructure	Achieved	The acceptance suite was extended to monitor the number of steps and tracks per event for each particle type. The statistical power of the comparison of the total visible energy was evaluated for samples of one to 20000 events. The study compared two alternatives for hadronic elastic scattering, demonstrating the resolving power of the chosen event sample of five thousand events to discern differences resulting in variations of total energy deposition of 4%. In addition a tool has been created for future studies of this type.
SIMU-?	31.03.06	Geant4 development release including new tool for overlap detection at geometry construction and extensions to QGS	Achieved	The March Geant4 development release (ref-03) included several fixes and a number of improvements. A new precise elastic process for protons is included, which approximates the cross-section and final states for projectile for p (energies T=100 MeV to 2.0 TeV), d and 4He targets (T=30 MeV to 900 GeV). A new method for overlap detection was included already in release 8.0 (December 2006). Contributions from other Geant4 collaborators included a prototype Python interface module, which interfaces to key Geant4 classes using Boost.

Summary of Progress

The main activity during this quarter has been the preparation of the software releases that are going to be used in the various data challenges and combined test runs of the LHC experiments during this year. About half of the functionality of SEAL has been completely migrated to ROOT and the experiments and the AA projects had made considerable effort in adapting their software to use the packages that have been migrated. A detailed plan has been started to be prepared for the migration of the second half of the functionality. During the quarter many different releases and software configurations has been produced to help the experiments for the preparation of their production releases. POOL products such as CORAL and COOL are coming with new functionality requested by the experiments. ROOT has made one development release including new functionality plus several bug fix releases, Geant4 has made various patch releases to support the experiments, and new versions of GENSER and GDML has been released. The HepMC package was installed in the LCG external area and is maintained by FNAL effort.

Concerning architecture and platforms. The AA software is being adapted for the AMD64 architecture and this is almost complete. Certification and preparation for the new Linux SLC4 has been made. Next releases will be made available for this new platform.

Outstanding Issues since Last Report				
Milestones Changes and Actions				
<p>Two simulation milestones, which were level 2 milestones for the simulation project, has been added to the list to better reflect the progress made by the project.</p> <p>New longer term milestones will need to be defined for ROOT and POOL projects. This will be done during the current quarter.</p>				
References and Hyperlinks				
Milestones for Next Quarter			Status	Comments
ROOT-9	30.06.06	First version of CINT running directly with the Reflex data structures as part of the ROOT June release	In progress	The time scale for this milestone has been very tight. We cannot meet this milestone, it would have to be postponed to the December release.
ROOT-10	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.
SIMU-8	30.06.06	New generator level production framework: beta release (from SIMU-2, due 15.12.05)	In progress	A lot of effort in the quarter, driven by the development of the new CMS generator level simulation framework. Prototype is now available. It seems realistic to start migration to LCG during the 2nd quarter.
SIMU-9	30.06.06	Investigation of correction for test-beam data for validation of stand-alone simulation engines (VD617)	In progress	Just started to evaluate the corrections for digitization and reconstruction effects for some observables of the ATLAS electromagnetic barrel calorimeter test-beam of 2002. Preliminary results of this investigation should be available in the second quarter of 2006.
New Proposed Milestones			Comments	
SPI-5	30.06.06	Provide the tools for the pre- and post-build procedures for the AA projects and externals through a web interface, such that it can be done or triggered by the project release managers.	Pre-build procedures are mainly the allocation of the AFS space (AFS volumes) for the release. The post-build procedures consists of generation of documentation, tarballs, etc.	
COOL-4	30.06.06	Support for multi-channel bulk insertion operations.	Functionality not yet implemented from COOL-1 and with lower priority.	
Comments and Additional Information				

QUARTERLY STATUS REPORT				
Project Name			Date	
Deployment Area			23 April 2006	
Report Period			Author Name	
2006 Q1			Ian Bird	
Milestones for the Quarter		Status	Comments	
GRID DEPLOYMENT				
GD-2	15.11.05	SC4 detailed plan agreed (following November GDB)	Completed: March 7 2006.	The SC4 workshop in Mumbai on 10-12 Feb gained agreement in many areas: data and storage management, service support and operations, and detailed explanations of the goals of the experiments for SC4 service phase. The conclusions were written up and finally agreed in the MB on March 7.
GD-4	15.01.06	Finalise the Baseline Services specification for the initial LHC service 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 services must be defined - for all sites (EGEE/OSG/NDGF sites)	Completed.	The Baseline Services WG discussions provided input to various other meetings and sub groups. In particular, the outstanding issues and priorities for the middleware were listed in the "Flavia" list and provided as input to the EGEE Technical Coordination Group as LCG priorities. Discussions on VO box issues was carried on in a separate GDB-mandated sub group; the storage management issues including those related to SRM versions are discussed in the storage management group, chaired by Maarten Litmaath and re-energised following the Mumbai discussions. Discussions on the VOMS roles and groups required by the experiments are discussed in the GDB, and also in the TCG.
GD-9	31.01.06	OSG and NDGF Baseline Service Plans Plan of the baseline services on the different grid implementations for SC4	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. Some of the expectations of the experiments were exposed in the SC4 workshop in Mumbai.
GD-10	31.01.06	Operations monitoring metrics implemented Covering the metrics and criteria for success specified in the MoU	Delayed until 30/4/06	A new version of SFT has been developed within which to implement the metrics. The majority of the metrics tests are available and under test. Simple extractions to Excel of initial measurements shown. Expect the first full version available end of April.
GD-8	31.01.06	Baseline services available for setup phase of SC4	In progress - and continuing.	LCG-2.7.0 was deployed at the end of January and together with FTS upgrades (contiuos process) provides the essential baseline services, with the exception of bulk job submission in the WMS expected with gLite-3.0. This latter has seen many major problems during certification and PPS testing; and a full deployment may not be ready by May.
GD-5	28.02.06	System and application tests for SC4 integrated in the Site Functional Test (SFT) framework	See GD-10.	The integration of the new tests is in progress, expected to be finished at the end of April.
GD-6	31.03.06	Service availability measurement system in place	See GD-5	This is really a duplication of GD-5.

MIDDLEWARE AND SOFTWARE DEPENDENCIES				
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. As discussed in Mumbai new versions of dCache with SRM2.1 will be deployed as they become ready and interoperability with other SRMs is shown. Will be monitored by the SRM group.
EXT-2	30.11.05	(EGEE/JRA1) consolidated VOMS support in place and operational	In progress	VOMS service for LCG is up and running in LCg-2.7 and gLite-3.0. How this is used by individual grid services is under discussion and implementation. This will be a continuing process. DPM in gLite-3.0 implements groups and roles; dCache implementation under way; Castor-2 implementation is not first priority. LCAS/LCMAPS available for batch jobs with the gLite-3.0 CE, but exactly what set of mapping should be used is not yet clear.
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. Significant critical bugs found during certification and pre-production; not yet all resolved; some workarounds but not yet clear how will behave under stress.
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	New CE in gLite-3.0 still has outstanding bugs. Not yet production-ready.
EXT-5	30.11.05	(RAL/CERN) CASTOR 2 SRM – basic version	Completed Jan 31.	SRM v1.1 in Castor2 is in production use.
EXT-6	31.01.06	(RAL/CERN) CASTOR 2 SRM for SC4	In progress.	See EXT-5. Development and deployment of v2.1 SRM agreed at Mumbai.
EXT-8	31.01.06	dCache SRM for SC4 ready for test by the experiments	In progress	See above. dCache in production for SC4; Development and deployment of v2.1 SRM agreed at Mumbai.
EXT-7	31.12.06	Implementation of VOMS roles and groups – placeholder - this can be defined only after final report of the baseline services group (end Nov 05)	In progress	Implementations in different services are in different states. For storage management better understood than for other services. SRM implementations exist (DPM), under way (dCache) or planned (Castor). For batch systems mapping to local priorities not understood yet.
Summary of Progress				
Service Challenges:				
<p>The re-run of SC3 took place in January-February and demonstrated sustained disk-disk transfers at an aggregate rate of close to 1GB/s; with rates to each of the Tier 1 sites at or exceeding their target rates. The limitation was found to be a network switch at CERN. The subsequent disk to tape tests successfully demonstrated rates to tape exceeding 50MB/s at each of the 6 sites participating. The rate was not smooth as the disk-disk rates, and several sites did not participate. The follow-up in SC4 will be important. Several issues in the File Transfer Services were uncovered and addressed. The Gridview monitor was also improved and is now used regularly to monitor progress. In March several pre-SC4 tests have been performed with a new Castor-2 instance at CERN dedicated to this. Initial tests to most of the Tier 1s have been run.</p>				

Mumbai workshop:

The SC4 preparation workshop took place in Mumbai on 10-12 February. There were detailed presentations and clarification of the experiments' goals for SC4; and significant progress was made towards agreeing the details of functionality that would be available in the middleware expected to be deployed by SC4. In particular the priorities for data management developments were agreed. A storage management sub-group, chaired by Maarten Litmaath was set up and is monitoring the progress towards developing and deploying the SRM v2.1 versions of the various storage managers. Agreement on how to provide the essential storage classes for the experiments was also agreed. It was clarified and agreed that v2.1 versions of SRM would not be in place for SC4 and would be deployed gradually as they appeared. Data management tools are thus needed that can manage both v1 and v2 implementations. These agreements were followed up and ratified in the management board of March 7.

Middleware releases and gLite-3.0:

The LCG-2.7.0 distribution was released at the end of January and deployed to the majority of the EGEE sites in February and March. An important patch was released to address problems in interfacing with some batch systems that had caused "black holes" for jobs in some cases. A 64-bit version of LCG-2.7 was available on March 3. This version of the middleware is the fallback solution for SC4, assuming that FTS, storage manager, and data management tool upgrades can be provided. This would be the mode in which SC3 was run last year.

The preparation of gLite-3.0 began at the end of January. This is intended as the first converged gLite and LCG release. However, progress has been very slow with many serious problems being uncovered.

The readiness for SC4 in June is uncertain. The basis of this version is LCG-2.7 and gLite-1.5. The latter was not released until the end of January; integration with LCG-2.7 began immediately and many issues of compatibility, dependencies, and different base versions were discovered. However, the real problems were that the WMS of gLite-1.5 did not work correctly with significant problems in bulk job submission, in the gLite CE, and proxy renewal for long jobs not working. Daily meetings of the certification team with developers were started and many issues were addressed.

By the end of the quarter the majority of problems have been resolved or it was decided that they were not showstoppers for this release. However, at the end of the quarter there are still major uncertainties in the behaviour under stress of bulk job submission.

A version of gLite-3.0 was passed to the PPS on March 3, deployed and made available for users by mid-March as planned. However, the unresolved problems meant that it was of limited use to the experiments. A very high job failure rate on the PPS was immediately observed, with many underlying issues: not only that of the middleware itself but problems in configuring and running

Interoperation:

Several of the OSG sites have been visible in the EGEE information system and accessible through the RBs. A third workshop with Nordugrid was held at the end of March to continue to progress towards interoperability. Progress was made in bringing together the information systems, and job submission from EGEE to ARC is possible. A 3-day workshop with the Japanese NAREGI project was held - important for LCG since the Japanese sites (Tokyo University and KEK) would eventually like to provide LCG services via NAREGI as local support.

Other points:

The different operations and service challenge weekly meetings have been merged to be more effective and to reduce the overlaps.

Outstanding Issues since Last Report

- 1) gLite-3.0 readiness;
This is discussed above.

Milestones Changes and Actions

References and Hyperlinks

[Mumbai workshop conclusions](#)

Milestones for Next Quarter			Status	Comments

Comments and Additional Information				
<p>Response to ATLAS comment on FTS instability in the PPS: this was due to the fact that the PPS set up was delayed due to the delay in producing the gLite-3.0 candidate release. In addition the responsiveness of the PPS sites in setting up services was very slow. However, the FTS has been in full production, the production instance being continually updated for the throughput phase of SC4.</p>				

QUARTERLY STATUS REPORT				
Project Name			Date	
ARDA				
Report Period			Author Name	
2006 Q1			Massimo Lamanna	
Milestones for the Quarter			Statu s	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)	Done	See in the summary of progress section
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)	Done	
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)	Done	
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)	Done	
Summary of Progress				
<p>During this quarter this activity had no specific new milestones since they are in preparation. The proposed milestones concentrates on a first set of activity and will be discussed in the MB soon. Some of the activities have no explicit milestone on it. The experiment support contains the "experiment Integration and support" activities and the integration activities. By their nature these actives have no fixed milestones. Also the AMGA activity work plan is a support activity (bug fixes, new platforms) and it will be monitored by the experiments and deployment via the TCG</p> <p>I propose to consider all the milestones as fulfilled since the activity continued to integrate the components under our responsibility in the experiments system.</p> <p>The grid system used for the analysis (prototype) system and for the productions is virtually the same and ARDA is giving its contribution in helping to improve the new components also in connection with the different experiments system. The distinction between the gLite installation is disappearing due to the convergence onto the coming gLite 3.0 (containing components from the production stack and selected services from gLite 1.5).</p> <p>Best examples are for the WMS: the activity in the ATLAS task force; for R-GMA: its usage in the CMS dashboard prototypes.</p> <p>In addition, wherever possible, the same tools are used for analysis and productions activities (e.g. LHCb GANGA submitting to the LHCb production system, DIRAC but also to the gLite system; the experiments dashboards collecting data from the production systems etc...).</p> <p>The rest of the activity is concentrating in integrating components in the experiments system and supporting/evolving existing ones.</p> <p>Finally wherever appropriate (ALICE, ATLAS and LHCb for their end-user analysis system) ARDA effort has been used to prepare documentation (notably the ALICE User Guide) and in tutorial efforts.</p> <p>ALICE: The evolution of the ALICE system continues towards its final deployment. The major good news on top of the integration activity, is that the ALICE Analysis User Guide was made public end of January. It includes instructions for 'aliensh', 'ROOT Grid plugin' + batch-style analysis code and examples. Since then a few tutorials for many of users (several tens) have been organized with good success.</p> <p>GANGA: version 3.1:</p> <ul style="list-style-type: none"> - The new Ganga GUI, completely re-written in PyQT. - Logical folder. Jobs can be grouped in logical folders for easier reference. A job can be in more than one logical folder (a la google mail) - Automatic Job splitting. - Several handlers are available (prototype for GAUDI, ATHENA etc...) - Note that GANGA is in use/being considered by several communities: Geant4, Biomed, NA48. <p>Also in this case, there is an important activity of tutorials and demo (for ATLAS and LHCb); in addition GANGA has been demoed at the EGEE User Forum and it will be part of the EGEE final review.</p> <p>AMGA: The metadata system developed in ARDA is now part of the gLite software stack. The system is made available for GANGA, for evaluation as bookkeeping system in LHCb plus other non-HEP users.</p>				

Experiment Dashboard:			
<p>The database for the CMS prototype is being rewritten (prototype exists). The new system runs as an Oracle service at CERN and has a new optimized and improved schema. All the data collection and monitor data handling within the dashboard has been redesigned in order to avoid the limitation of the first version (especially triggered by incomplete monitor information). The new CMS dashboard is capable to collect data also from PheDex. We expect the new system to be in production in April.</p> <p>In parallel, an ATLAS prototype has been set up, using the new database. This system is being harmonized with the ATLAS infrastructure (in particular, it allows keeping under control the LCG production). We have started to connect the system with the new ATLAS data management system.</p>			
Outstanding Issues since Last Report			
Milestones Changes and Actions			
<p>New programme of work discussed with the experiments. There is an agreement on the content, which will be presented to the LCG MB. The detailed timetable (with milestones and decisions points) is in preparation.</p>			
References and Hyperlinks			
<p>EGEE User Forum: http://egee-intranet.web.cern.ch/egee-intranet/User-Forum/index.html</p>			
Milestones for Next Quarter			Status
GANGA-AT1	Jun-06	Continue the evolution of the system within the experiment framework. The tutorial activity should continue and a first integration with the experiment data management should be available.	
GANGA-AT2	Nov-06	Production version exposed to the user community and used in SC4.	Report on analysis activity: # Active users # Feedback # Comparison with other tools in use in the collaboration
GANGA-LH1	Jun-06	Continue the evolution of the system within the experiment framework. The tutorial activity should continue. Splitting functionality should be fully integrated.	
GANGA-LH1	Nov-06	Production version exposed to the user community and used in SC4.	Report on analysis activity: # Active users # Feedback # Comparison with other tools in use in the collaboration
DASHB-AT1	Jun-06	Beta version to demonstrate aggregation of data from a variety of sources, including views of the ATLAS production system	
DASHB-AT2	Nov-06	Version 1.0 in use in SC4	Improved version in use in SC4: # Management view # Detailed views as defined by the experiment # Drill-down functionality to track problems # First automatic error detection
DASHB-CM1	Jun-06	First full release, including CMS specific services information. First functionality of troubleshooting part (access to error logs, correlation of the error information)	The data used by the dashboard should enable to produce time series of meaningful quantities (e.g. error type/#jobs vs. time) to evaluate the service and its usage
DASHB-CM2	Nov-06	Version 1.0 in use in SC4	Improved version in use in SC4: # Management view # Detailed views as defined by the experiment # Drill-down functionality to track problems # First automatic error detection
Comments and Additional Information			

QUARTERLY STATUS REPORT				
Project Name			Date	
Distributed Database Deployment			11/4/06	
Report Period			Author Name	
2006 Q1			Dirk Duellmann	
Milestones for the 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.	Done.	The database hardware has been acquired and installed at all sites. Acceptance tests have been concluded and the database servers are available for the pre-production phase.
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	The experiment work on software tests of their database application has progressed. The definition of payload data for detector conditions is still in progress and is expected to still evolve over the next month.
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.	Done.	A database workshop was held at CERN where the sites and experiments reported on their progress.
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)	in progress	COOL and Frontier installations exist at CERN and several Tier 1 sites. The completion of installations at all sites is expected to conclude during April.
DBS-6	31.03.06	Tier-1 services starts - milestone for early production Tier 1 sites	Done.	Level 1 - see above
Summary of Progress				
<p>During the last quarter the database and frontier setup for the prep-production phase has been largely completed. Progress was monitored during the now weekly 3D phone meetings and two database workshops (one held in Feb at CERN, one dedicated database administrator workshop hosted by Rutherford Lab). All tier 1 site have made 2-3 node database clusters available for the next phase and also the tier 0 setup - an increased of DB cluster capacity by factor two and a redundant 3 node Frontier production system - is now setup and ready for larger scale deployment. The experiments have continued to refine their database models for condition data and all software frameworks integrate now the required conditions s/w, but concrete condition data models for several important subdetectors are still missing.</p>				
Outstanding Issues since Last Report				
<p>CMS has proposed that 3D should now take over the follow-up of the SQUID installations at the LCG sites. This will be done in a combination of 3D meetings and LCG service meetings as the responsibility for SQUID is on some sites with the DBA teams, on other sites with the Sysadmin teams.</p> <p>Several sites noted that the service review workshop (milestone DBS-7) may in fact come too late during their budgeting or acquisition cycles for any larger change in their database capacity for the October service.</p>				

Milestones Changes and Actions			
<p>As the experiments are not in the position to determine a target throughput rate for their replication activities we propose to put a milestone on the project, experiments and sites to perform a larger scale replication ramp up test (with all pre-production sites) by the end of May. This milestone should deliver throughput rates with artificial conditions data work load which can be achieved and sustained over at least a day with the current production setup. These rate would then be used as input to the experiments discussion on their calibration models.</p> <p><i>We also propose to add two project documentation milestones:</i> <i>end of May - Replication Technology Writeup (summary of replication test with Oracle Streams and Frontier)</i> <i>end of August - Backup/Recovery Strategy Writeup (summary of experiment requirements and site preparations for database recovery)</i></p>			
References and Hyperlinks			
<p>Database Readiness Workshop @ CERN Agenda: http://agenda.cern.ch/fullAgenda.php?ida=a058495 Database Administrator Workshop @ RAL Agenda: http://agenda.cern.ch/fullAgenda.php?ida=a061192</p>			
Milestones for Next Quarter		Status	Comments
DBS-7	31.05.06		
		<p>Service review workshop >> hardware defined for full production. Experiment and site reports after first 3 month of service deployment. Define db resource requirements for full service. Milestone for experiments and all tier 1 sites.</p>	
DBS-9	31.05.06		
		<p>Database replication throughput test Perform a systematic throughput test on the now available production setup. In particular measure the sustained replication rate which can be achieved between online and offline setups and between CERN and the tier 1 sites which participate in the preproduction phase.</p>	
DBS-10	31.05.06		
		<p>Replication Technology Write-up Produce a document summarizing the experience gained with streams and frontier/squid in the 3D testbed and outline the expected advantages and issues with both approaches to guide the deployment test with the production setup.</p>	
DBS-11	31.05.08		
		<p>Backup/Recovery Strategy Write-up Produce a document summarizing the required service level for database backups at tier 0, 1 and 2 based on the experiment input. This document should include the main recovery scenarios and describe the recovery procedures and expected latencies.</p>	
DBS-8	30.09.06		Level 1
		<p>Full LCG database service in place - milestone for all Tier 1 sites</p>	
Comments and Additional Information			

QUARTERLY STATUS REPORT				
Project Name			Date	
Grid Deployment Board			24.04.06	
Report Period			Author Name	
2006 Q1			Kors Bos	
Milestones for the Quarter			Status	Comments
GDB-1	11/1/06	GDB meeting at CERN	done	Meeting summary on the GDB wiki
GDB-2	8/2/06	GDB meeting at CERN	done	Meeting summary on the GDB wiki
GDB-3	8/3/06	GDB meeting at CERN	done	Meeting summary on the GDB wiki
GDB-4	5/4/06	GDB meeting at CASPUR in Rome	done	Meeting summary on the GDB wiki
Summary of Progress				
<p>Summary of the meeting and conclusions and decisions are in the Summary report. Officially approved detailed minutes are also available.</p> <p>See the GDB wiki: https://twiki.cern.ch/twiki/bin/view/LCG/GridDeploymentBoard</p>				
Outstanding Issues since Last Report				
Action items can be found on the GDB wiki				
Milestones Changes and Actions				
No GDB meeting in May 2006				
References and Hyperlinks				
GDB wiki: https://twiki.cern.ch/twiki/bin/view/LCG/GridDeploymentBoard				
Milestones for Next Quarter			Status	Comments
GDB-5	7/6/06	GDB meeting at CERN	open	agenda at https://twiki.cern.ch/twiki/bin/view/LCG/GridDeploymentBoard
GDB-6	5/7/06	GDB meeting at CERN	open	agenda at https://twiki.cern.ch/twiki/bin/view/LCG/GridDeploymentBoard
Comments and Additional Information				



QUARTERLY STATUS REPORT				
Project Name			Date	
ALICE			10.05.2006	
Report Period			Author	
2006 Q1			Yvez Schutz	
Milestones for the Quarter			Status	Comments
MS5	Sep-05	Start of distributed analysis	Done	An advanced prototype will be released end of January and tested by selected users. 11/04/06: The proptotype has been released end of January. A few users are exercising the analysis on a reduced scale. The software is now being debugged and improved following the users feedback.
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. 11/04/06: No progress.
MS8	Dec-05	Preliminary implementation of algorithms for alignment and calibration ready for all detectors	Done	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. 11/04/06: All detectors have implemented the alignment and calibration procedures within the general framework.
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	Done	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. 11/04/06: The AliRoot version to be used for the DC06 has been released. It includes all the functionalities foreseen to be tested with the DC06.
MS10	Jan-06	Start of PDC06	Done	Delayed until March 2006. 11/04/06: The production phase of PDC06 has started on limited sites. More sites will be integrated progressively starting with the T1.
MS100	Jan-06	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	Done	11/04/06: We have participated to the transfers and done some tests with FTS
MS101	Feb-06	Start of distributed analysis evaluation by selected users	Done	11/04/06: A tutorial on the distributed analysis has taken place in March (30 participants) and an other one will take place mid-April. The evaluation by these new users is now slowly taking off.
MS102	Mar-06	T0-T1 “loop-back” tests at twice the nominal rate (CERN). We run our bulk production and send data back to CERN	Delayed	11/04/06: Delayed to mid-April, with the start of PDC06.
MS103	Mar-06	PROOF ready at CAF	Done	11/04/06: In discussion with CERN resources manager in the context of SC4. A limited cluster (30 WN), properly configured will be made available for the validation of PROOF. 08/05/06: PROOF has been installed on a cluster of 40 WN.
Summary of Progress				
<p>The AliRoot code is reaching completeness with the final implementation of the raw data format, calibration procedures and alignment procedures using the ROOT geometric modeller. Detailed studies on the impact of misaligned geometries on the tracking performances in the central tracking systems and the MUON arm are well advanced. The overall calibration procedures including online calculation of calibration parameters are examined detector by detector. The program, called SHUTTLE, for collecting online parameters to make them available in the offline Calibration Data Base has been implemented. A complete documentation of AliRoot including a primer for new users has been released.</p>				

The distributed computing environment has reached sufficient maturity to make it available to general users. It incorporates the services distributed through the LCG middleware and ALICE specific services not available elsewhere. VO boxes hosting the ALICE services have been installed in all Tier1 sites and all Tier2 sites pledging resources to ALICE. The production phase of the Physics Data Challenge 06 will start on 10 April with the simulation of pp events defined by the Physics Board. Although all the utilities exist, distributed analysis has not yet been tested in a realistic environment and at a realistic scale. This testing will be a major challenge until the end of the year. A complete documentation and primer of the usage of ALICE distributed computing environment and distributed analysis has been released.

Distributed analysis based on the Parallel ROOT Facility (PROOF) and on the xrootd I/O subsystem will be the baseline for performing calibrations and prompt analysis on the CERN Analysis Facility during data taking and could be adopted by the Tier2 sites dedicated to mainly analysis. A prototype system will be deployed on medium-sized farms at CERN, CCIN2P3 and FZK to be exercised by general users.

Outstanding Issues since Last Report

The computing resources pledged so far by Funding Agencies for ALICE are well below the needs documented for and validated by the Computing Resources Review. New developments in the EGEE/LCG middleware services are planned to be released as late as October 2006. The adoption of new services or implementations at such a late time could have negative impacts on the availability of stable software in due time for the first LHC physics runs. The personnel situation for offline computing remains critical, stabilisation of key personnel is mandatory to ensure availability and quality of services to users.

Milestones Changes and Actions

Most of the coming milestones are depending on the time scale of SC4. The software part of PROOF is ready for validation, the IO/xrootd part still need some development, the start of the exercise is pending on the availability of a properly configured cluster at CERN. The validation of the distributed analysis on a large scale (many users, distributed data) is still an issue.

References and Hyperlinks

Milestones for Next Quarter		Status	Comments
MS104	Apr-06	T0-T1 disk-disk (nominal rates) disk-tape (50-75MB/s). First chance to push out data, reconstruction at CERN	11/04/06: Coupled to MS100. Delay in line with the completion of the disk-to-disk transfer T0->T1
MS105	Apr-06	First tests with PROOF at CAF	11/04/06: this milestones will be delayed pending on the decision of the CERN resources management to make a properly configured cluster available for PROOF. 08/05/06: The PROOF cluster has been configured at CERN. Tests will start soon.
MS106	Jul-06	T0-T1 disk-tape (nominal rates). Second chance to push out the data. Reconstruction at CERN and remote	11/04/06: On schedule, in coordination with SC4. Based on the resources availability at T1 centres (networks and storage), the transfer rates may be revised
MS107	Jul-06	T1-T1, T1-T2, T2-T1 and other rates according to the rated defined in the TDR.	11/04/06: On schedule, in coordination with SC4. T1 support for T2 centres (without announced regional support) is being negotiated.
MS108	Sep-06	Scheduled analysis challenge in T1s.	11/04/06: There is no delay anticipated, pending however on the successful completion of phase 1 of PDC06 (see MS10)
MS109	Sep-06	Unscheduled challenge in T2s.	11/04/06: There is no delay anticipated, pending however on the successful completion of phase 2 of PDC06 (see MS108)

Comments and Additional Information

QUARTERLY STATUS REPORT				
Project Name			Date	
ATLAS			10.05.2006	
Report Period			Authors	
2005 Q4			Dario Barberis	
Milestones for the Quarter			Status	Comments
ATL-CSC-05-01	11/30/2005	Start of simulation production for Computing System Commissioning.	Delayed to May 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 (May 2006). Depends on milestone ATL-REL-06-01.
ATL-REL-06-01	1/31/2006	Software release 12: production release for Computing System Commissioning and early cosmic-ray studies; completion of the implementation of the Event Data Model for reconstruction.	Delayed to May 2006.	See above comments.
ATL-CSC-06-01	2/28/2006	Start of Computing System Commissioning.	CSC started at the end of March 2006.	Now coordinated with SC4 activities. Tier-0 tests and production operations that do not need SC4 infrastructure will take place starting in March 2006.
Summary of Progress				
<p>Integration of new software infrastructure components that are necessary for Computing System Commissioning operations (mainly ROOT5 for schema evolution, COOL 1.3 for the Calibration Data Challenge) took much longer than expected. The implementation of the realistic detector geometry (including miscalibrations and misalignments) is also in progress. Release 12, including all these features, will be used for the large-scale CSC simulation production from May onwards. Tier-0 throughput tests were run successfully in January as last part of the ATLAS SC3 activities; nominal rates were reached and held for >24 hours. Preparations for SC4 tests (Tier-0 and data distribution) later in April are making progress.</p>				
Outstanding Issues since Last Report				
<p>The functional tests of data distribution that we had foreseen as preparation for SC4 starting in late March could not succeed as the required middleware components and infrastructure (FTS on the pre-production service) were not stable enough for us to work with them. This lack of preliminary testing will generate a delay later on during the year as a larger part of our software will have to be tested at the same time within a limited time window.</p>				
Milestones Changes and Actions				
References and Hyperlinks				

Milestones for Next Quarter			Status	Comments
ATL-CSC-05-01	11/30/2005	Start of simulation production for Computing System Commissioning.	Delayed to May 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 (May 2006). Depends on milestone ATL-REL-06-01.
ATL-REL-06-01	1/31/2006	Software release 12: production release for Computing System Commissioning and early cosmic-ray studies; completion of the implementation of the Event Data Model for reconstruction.	Delayed to May 2006.	See above comments.
ATL-CSC-06-02	4/1/2006	Integration of SC4 components with ATLAS software stack.	In progress	
ATL-CSC-06-03	5/15/2006	Definition of SC4 detailed milestones.	Planning activity in progress.	
Comments and Additional Information				

QUARTERLY STATUS REPORT				
Project Name			Date	
CMS			10.05.2006	
Report Period			Author Name	
2006 Q1			Lothar A. T. Bauerdick	
Milestones for the Quarter			Status	Comments
CPT-212/C	Mar-06	Computing Systems Ready for Cosmic Challenge	completed	Data transfer from site 5 to remote sites demonstrated with few hours latency by end of April by C4 Task Force led by T. Wildish
CPT-402/C	Apr-06	Computing Systems Ready for SC4	delayed	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 403/C-1	Mar-06	Computing Systems ready for start of simulation production w/ new EDM/FW	completed	initial prototype of Mcprod system completed, 100GB of data produced. First version released to operation by end of April. Led by P.Elmer
Summary of Progress				
<p>Details of interactions between new CMS Software Framework, Data Management and Production Workflow are being defined. A prortype now exists for DataManagement catalogs and MonteCarlo Production Tool. User's analysis tool (CRAB) is being ported to new Software Framework. Data Trasnfer component (PhEDEx) is being modified to use FTS as underlying file transfer tool. CMS Tier0 project has started. Job Submission Robot to generate analysis load during SC4 operational.</p>				
Outstanding Issues since Last Report				
<p>SC4 middleware (gLite 3.0 stack) still not usable on preproduction system. SRM and CERN-Castor2 instabilities during completion of Physics TDR MonteCarlo production caused significant operational load and delayed setup of tools for SC4 data transfer activities</p>				
Milestones Changes and Actions				
<p>CPT 402/C will be completed by end of May in time to rerun SC3 service phase in June as part of SC4. Small scale test of all functionalities needed for CSA06.</p>				
References and Hyperlinks				

Milestones for Next Quarter			Status	Comments
CPT-402/C	May-06	Computing Systems Ready for SC4		delayed to end of May since SC4 starts June 1st
Comments and Additional Information				

QUARTERLY STATUS REPORT				
Project Name			Date	
LHCb			10.04.2006	
Report Period			Author Name	
2006 Q1			Nick Brook	
Milestones for the Quarter			Status	Comments
Mar-06	Final alignment Strategy		Ongoing	Expected to be finished end of April
Summary of Progress				
<p>The VELO produced a report detailing their alignment procedure; the other tracking detectors are now also studying their alignment strategies. Work has commenced understanding how to perform the global alignment; a discussion document is expected in April.</p> <p>The LHCb Workload and Data Management System (DIRAC) was reviewed in readiness for the 2006 Data Challenge. The tools are being finalised to allow automated processing to be triggered as data files become available.</p> <p>Core developments in Gaudi have made use of the developments in the LCG application area following the merging of ROOT and SEAL. We have retired the use of the CLHEP library from the LHCb-specific parts of the applications (though it is still necessary for GEANT4) and LHCb has contributed to testing and improving the new ROOT geometry and Linear Algebra packages.</p>				
Outstanding Issues since Last Report				
Failed to meet metric for T1-T1 transfers				
Milestones Changes and Actions				
References and Hyperlinks				
http://savannah.cern.ch/projects/dirac2/ https://twiki.cern.ch/twiki/pub/LHCb/DiracReview/Reviewreport.doc				
Milestones for Next Quarter			Status	Comments
31.05.06	Start data processing phase of DC'06 (over a 2 month period): (I) data distribution (125 TB) (II) reconstruction/stripping (CPU: ~300kSI2k.monnth; MSS:~200TB) (III) distribution of stripped data (2.2 TB disk storage at each Tier-1)			planned to start soon after gLite3 deployed in production environment
Comments and Additional Information				