U.S. CMS Software and Computing

Progress Report for the 4th Quarter FY2005

Technical, financial and management status is reported for the period of June 30th to September 30th , 2005.

Technical Status

U.S. CMS Software and Computing efforts were driven by the major milestones and activities for FY05:

Reliable and automated transfers of data between the CERN Tier-0 and USCMS Tier-1 centers
Participation in the LCG service challenges to commission a worldwide Grid service for CMS
Preparation, development and integration of CMS software and computing systems

The computing technical baseline of the CMS experiment was documented in the Computing TDR, that was submitted to the LHCC for review in June, and that was being reviewed in early October 2005. USCMS was heavily involved in developing the concepts and writing the C-TDR. With P.Elmer of Princeton U. taking the co-leadership of the Technical Program area and Ian Fisk of Fermilab taking co-leadership of the Integration Program of CPT-Computing the US is continuing to strengthen its engagement in overall leadership of the CMS computing project.

USCMS Tier-1 Facilities

The USCMS Tier1 facilities group finished procuring their FY05 computing equipment. This included:

Worker nodes: As previously reported, USCMS tried to limit its financial exposure to problems of defective equipment in the procurement process. Instead of placing the FY05 order for all nodes at once, a procedure was worked out to accepting only a few racks of nodes at a time from the vendor.
We accepted the 1st rack (40 nodes) and had it in production within 2.5 weeks as previously reported.
Much was learned during the delivery of the 1st rack, such as

(a) optimizing logistics within lab to ensure proper accounting happens quickly,
(b) getting IP addresses assigned before the nodes arrived so tests can be immediately started,
(c) installing the condor batch system to the automate the benchmarking tests.

The second set of nodes received from the vendor KOI (racks 2, 3, 4 – 120 nodes) arrived during this period . All nodes were installed and ready for testing within a few days. Unfortunately, the results of this testing was not as good as the first rack and many nodes were crashing and failing the integration tests. The vendor promptly responded and performed proper corrective actions for new nodes. For example, for nodes that were failing the 2-week burn-in tests, these type of actions were taken: CPUs that were possibly jarred during shipping were reseated in their sockets, the motherboard BIOS was reset to default values, and disks were replaced. The infant failure rate stopped after these actions, but failures still occurred. The

vendor then told us that they used a different brand of memory than the 1st set, memory that was still acceptable according to the bid and the motherboard vendor. They offered to replace the new memory for the type on the first rack of nodes on a few nodes that were failing the burn-in tests every day as a test. They did this, and the failures stopped.
We then had a several day delay while KOI ordered new memory for all the nodes and while they came out to replace all memory on the new worker nodes. At this point we reset the burn-in clock and started the tests over. This time, the 120 new nodes passed the burn-in test. We immediately put them into production service and they are being used heavily. Although, the process took us one month instead of the three week goal, we are still pleased that bad hardware was caught during the integration testing minimizing potential impact on CMS users.
After the 120 new nodes were installed, we decommissioned our oldest worker nodes from production and are currently configuring them for test and development clusters.
Finally, we asked for delivery of the last set of worker node options (three more racks, 120 more nodes) from KOI and delivery is scheduled for late October. For next year's order, we plan on adding a clause that requires the winning vendor to state what has changed in new shipments compared to the last one, even if the parts are still within the bid spec.

Data disk: The OS of servers attached to the data disk was upgraded to Scientific Linux 4. This OS and an upgraded java have allowed us to max out the GigabitEthernet network links (100 MByte/s) on the data disks serving the dCache pools. We are just starting to see what rate we can achieve by aggregating 2 GE lines to the servers. The ~30 TB of Nexsan Atabeast Storage Arrays that were ordered during the last reporting period have arrived and we are putting them into production service now. We have also freed up one older Infortrend storage array that we were using for NFS service and are putting this into production dCache service. This will bring USCMS resources to the desired ~100 TB of dCache data disk. A big challenge for FY06 is ordering and commissioning 600 TB of dCache data disk. New models or storage arrays are out for Nexsan and Infortrend and we plan to evaluate both for the next purchases, as well as any other models that can compete with them.

User disk: We added an Infortrend storage array, also previously used for NFS service, into the IBRIX global file system. This file system is 'controlled' completely by the physics group conveners and has been very popular with CMS users because it does not have quota restrictions. We are working with IBRIX to be able to provide a quick list of who is using space on this files system, so the physics conveners have this information available when space limitations start to happen. The Infortrend/Nexsan disk is substantially cheaper than the higher quality STK/LSI disk we use for the quota-enabled/managed IBRIX file system and therefore it may be easier to extend this file system than the higher-quality LSI disk. Current tests indicate roughly identical performance on worker client nodes between the Infortrend disk and LSI disk, although we expect the LSI disk to perform better when overall rates increase. One problem has developed with IBRIX that affects jobs submitted from the grid – Globus uses a hard-link counter as a lock for job concurrency. This does not work reliably in IBRIX, and the IBRIX engineers have not been able to determine why. We have a work-around in place, where we use a symlink in IBRIX to point to a local directory. Globus then creates the hard links in a local file system thereby avoiding the problem. A search on the web indicates this identical problem also happens with NFS file systems, and possibly other global file systems. We are planning a major change in client access soon – switching from UDP to TCP in the native client connections. IBRIX engineers indicate the change causes a 5% performance degradation, but solves many problems associated with dropped network packets.

User Analysis Facility: Two major changes happened during the last reporting period. First, we eliminated the old FBS batch environment, and now require users to submit to the general condor batch queue. We announced this many times to our users and the transition went well. Removing the FBS environment made the UAF a completely interactive cluster. This allowed us to reduce the number of nodes and return them to the general condor worker node pool where they are available and used by all condor job submissions. Second, we eliminated the old FBS interactive load balancing method and now require users to connect interactively to the UAF nodes using the native hardware load balancer of the CISCO 6509 switch. This was also announced many times, but didn't go as smoothly as the first item. The trouble was that it required clients to upgrade the software on their nodes to more modern versions of ssh. Older versions of the client software had errors in them, errors that were previously ignored, but do now matter. For the users that for whatever reason could not upgrade, we provided several direct access nodes for their UAF interactive work.

Data transfers: Work continued to automate and optimize data transfers within USCMS nodes at FNAL and also between the Tier1 site at FNAL and the 7 USCMS Tier2 sites. The key improvement during this period was the introduction of the multi-IO queues in the dCache software. This allowed us to have separate quotas for LAN and WAN data transfers. This is important because of the different nature of the CMS LAN and WAN transfers. For LAN transfers, CMS software wants to open many files, but typically only reads a few bytes at a time from each file. Therefore this LAN mode wants very high transfer quotas and we typically reliably operate with more than 300 transfers per dCache pool. WAN transfers, by their very nature, are always high-rate full file transfers. Gridftp further optimizes WAN traffic by opening 10 to 20 parallel streams to overcome bandwidth-delay problems. This means each WAN transfer is actually ~20 high-rate transfers, and we have found that each dCache pool can typically handle ~3 WAN transfers at a time. With the introduction of the multi-IO queue feature in dCache, these two competing transfer modes can now have separate quotas and co-exist on the same pools. This has stabilized operations and increased data traffic. We recently achieved a record day – 200 TB/day, with 1/3 of that coming from the main dCache pools and 2/3 coming from the resilient dCache pools deployed on the worker nodes. This included WAN, LAN, and Enstore tape restores and stores – that is, the full set of IO modes required for USCMS during real data taking.

1.2 USCMS Tier-2 Facilities Program

The US-CMS Tier-2 program has begun its evolution from a research and development project to a production system, with important steps towards deployment of a full system. This has been achieved by starting the acquisition of the necessary computing hardware, hiring the needed staff at the Tier-2 sites, and setting (and achieving) the goal of full participation of CMS' Service Challenge 3 (SC3).

Six of the seven sites (Caltech, Florida, Purdue, Nebraska, UC San Diego and Wisconsin) received their first allocation of money early in this quarter, and for reasons of availability of funding, MIT is starting in September. All quickly turned their attention to the goal of purchasing the hardware needed for "20%" of the system required at LHC startup. This milestone was established at the Tier-2 workshop in May, and "20%" was defined to mean 60-80 dual CPU worker nodes and about 20 TB of disk-based storage deployed as a dCache storage element. By the end of the quarter, all sites had decided what hardware they wanted to purchase, and had placed their orders, with most also taking delivery and deploying. All were able to meet the 20% milestone with the available funding. The funding for this is a combination of USCMS research program and DISUN funding.

In addition, funding has been provided for two FTEs at each site to operate and maintain these systems. Thus, there are a total of fourteen positions, and by the end of the quarter, 10.7 FTEs had been identified and hired for the project. This is an important step towards establishing regular and stable operations of Tier-2 sites well in advance of LHC startup.

With these resources in place, the sites turned their attention towards deploying the necessary services. The goal was to be ready for the service phase of SC3, which began on September 15. This service challenge was meant to test the basic functionality of the entire CMS computing system, which at Tier 2 meant handling data transfers from the Tier-1 site at FNAL, processing analysis jobs submitted through Grid interfaces, and generating production simulation samples — simultaneously. To be ready for this, the sites had to deploy a great variety of computing services, such as the dCache storage system, the Open Science Grid software stack, the databases needed to keep track of hosted datasets, and of course the CMS software itself. Through the hard effort of the staff at each site, all but MIT (which by agreement had a late start) were prepared for SC3 in time for the start, and have been integrated into the challenge as needed.

This demonstrates that the sites are on track to perform their assigned tasks at LHC startup.

1.3 Grid Services and Interface with the Open Science Grid (OSG)

The Open Science Grid software release 0.2 was completed and deployed, resulting in the "ribbon cutting" for the Consortium at the collaboration meeting in Milwaukee July 20th. The scale of the OSG is 3x that of Grid3 with more than 14,000 CPUs accessible through the common infrastructure. The US CMS S&C contributions included integration testing, documentation, coordination. Burt Holzman completed the procedures and organization for the US CMS support of OSG working closely with the Fermilab Computing Division Helpdesk and Indiana Grid Operations Center (iGOC). Support contacts from the Tier-2 sites are included also. Tickets can now flow automatically from Fermilab to the iGOC and support tickets are flowing from OSG (and LCG) to US CMS operations.

Scripts were written to automatically install, configure and test new versions of the Open Science Grid software cache and are extending it to provide a framework for similar tests of the US CMS software that will be installed and run on OSG. These scripts have already found several problems in the provisioning.

Grid security and authorization, and Grid accounting, ramped up their contributions. We made sure the components of the Authorization infrastructure work with the new version of the Virtual Data Toolkit that includes the GT4 GRAM job interface. We worked on the Grid accounting project. The requirements and design phase were completed, and signed off by the OSG technical group (Monitoring and Information) and deployment activity, which are overseeing the activity. Initial Tier-1 accounting information was readied for the LCG and is currently undergoing internal review.

The integration and deployment of VOMRS for the LCG entered the user testing phase. The software was interfaced to new versions of the Virtual Organization Management Software (VOMS) delivered from the EGEE gLITE project. ORACLE is now being supported as an additional backend database.

The US CMS PPDG and iVDGL grid project effort contributions to Grid Services and interfaces area (GSI) for the last quarter include: extensions for OSG of the Grid Catalog and support for Clarens use by the DCT area; integration of the authorization service for the Storage Element into the SRM/dCache infrastructrure; and deployment of Storage Elements at the iVDGL Tier-2s. The work in the Condor Project has resulted in Condor-C is now being part of the gLITE releases. Tests started at the US CMS Tier-2s for its use in the US.

Progress was made on interoperability with the LCG with the deployment of common Information Providers which are used to advertise and select resources for jobs submitted to the LCG.

1.4 Application Services

USCMS engineers are working with the international CMS efforts in the CPT-Computing project to provide the application-level computing services. The requirements, interfaces and functionalities of the application areas are now reasonably well described in the technical baseline documented in the CMS Computing TDR.

The initial design of the Datset Bookkeeping System DBS proceeded in this quarter, with a design report being released within the CMS data management group. People at Fermilab worked on a prototype implementation of the database server, and on the mechanisms to fill the DBS schema from the pre-existing refdb (MC production database) information.

Unfortunately CMS lost one of its lead developers in this area, with G.Graham leaving Fermilab. It took until September that he could be replaced by moving an experienced developer and designer with a lot of expertise on data management gained in the SAM project, S.Viseli, into the CMS project. However, sepcific deliverables in the data management area like the prototype of the DBS was slipping by several months. One of the US engineers based at CERN, L.Tuura, took up coordinating the DBS efforts in particular to make sure that the required deliverables for the Service Challenge 3 were delivered.

The PhEDEx data transfer system was deployed at all SC3 sites and monitoring and operations support tools were developed and installed. PhEDEx orchestrates the production dataflows of moving datasets to the hosting sites. Also PhEDEx is being successfully used to orchestrate the SC3 related data transfers.

Other efforts included making progress on the definition of the data flow from the HLT into the Tier-0 center and the related data formatting and data bookkeeping issues.

In the September PMG meeting project oversight agreed that P.Elmer would take up coordination of this area and since then he has become part of the Project Execution Team and is closely coordinating these efforts.

1.5 Distributed Computing Tools

USCMS continues to fulfill its obligations for generating Monte Carlo for CMS. Within the last year the two production centers, USMOP based at Florida (submitting CMS jobs to Grid3 and OSG resources), and GLOW at Wisconsin have each generated more than 40 Million events. This has to be compared with roughly 100 Million events generated at CERN, and about 300 Million generated globally during the same time frame.

The total MC produced in the US is neither human nor hardware resource limited. It is limited today only by the number of events the CERN based production team assigns to the US. CERN is effectively providing an organizational bottleneck between requests for MC from US based physicists and the two US production teams. We have started to address this towards the end of Q4, and first results are already visible.

ApMon and JobMon were integrated into the ShREEK runtime workflow management system. The former is used to send out status messages from running jobs to the CMS dashboard. The latter provides appropriately privileged users read-only access to the runtime execution environment for debugging purposes.

Given the intricacies of deploying the CMS software on diverse grid sites across OSG, a common interface layer for grid jobs was developed called CMSSoftDB. This is a lightweight python interface to an XML DB structure used to contain site and CMS software version specific details, and present a uniform interface to the User. Development of CMSSoftDB was stimulated by the collaboration with DISUN on CMS software installation and validation. MCPS is now used as validation tool as described below. The content of CMSSoftDB is maintained as part of the software installation effort described below.

Initial design discussions were started for a complete redesign of the Monte Carlo production system for CMS in order to fully benefit from the redesign of the Event Data Model and framework, as well as the various data catalogues and provenance system. We expect the development of the new MC production system to be the main focus for the next quarter.

DISUN contributed to the USCMS DCT area in Q4 in three broad areas: development of CMS specific tools, development of grid services, and integration effort focused towards making sure the OSG 0.4 release, planned for late 2005, satisfies CMS needs.

Development of CMS specific tools has focused on CMS software release installation and validation. DISUN has taken on the responsibility for CMS software release deployment, and validation for the Open Science Grid. Tools to do so were designed, developed, and deployed within Q4. After an initially painful integration period, we are now in a position to deploy and validate a new CMS release on all USCMS Tier-2 centers within a day. One of the benefits of this procedure is rapid feedback to the cmsi team responsible for providing CMS releases as rpm's for deployment within global CMS.

The software installed is used by CRAB as part of service challenge 3, and planned to be used by the new MC production workflow. DISUN contributed the automatic installation tools, and integrated MCPS and CMSSoftDB, both cms-dct products, for the purpose of site configuration and validation. Advertisement of the installed releases is done via the GIP.

At present, USCMS Tier-2 sites are not accessible by the standard CRAB version in CMS because our sites are not compliant with the resource broker (RB) that CRAB uses. In Q4, we started to address this in two ways. A modified CRAB version was created that uses Condor-G submissions to the sites directly instead of the RB. We expect these modifications to be fed back into CRAB in the next quarter. This is a joint effort between DISUN and Fermilab. In parallel, we have started to work with the USCMS Tier-2's on RB compatibility. USCMS Tier-2 site validation thus comprises of three steps. First software installation is validated using MCPS, followed by PubDB validation using the modified CRAB version, followed by RB compatibility validation, again using CRAB. At this point all 6 operational sites have passed steps one and two, but only two sites have passed step three.

In addition, a first prototype of the MCPS web service was created within the Clarens Web Services framework, and demonstrations where given at the OSG meeting at Milwaukee, as well as iGrid 2005, San Diego. This prototype is the first implementation of the "me - my friends - the grid" paradigm described e.g. in the CMS Computing TDR. All CMS specific persistent services are deployed at Tier-2 centers ("my friends") thus minimizing the requirements on interfaces installed by the user or the grid. The Clarens interface to MCPS distinguishes user and administrator functionality, and presents these to a user based on the extended proxy presented by the user to the service. This work was done in close collaboration between DISUN and Fermilab.

In the area of grid services, DISUN is contributing to work on Condor-C and the Edge Services Framework (ESF). Both of these are being worked on in close collaboration with OSG, EGEE, and LCG. Within Q4 DISUN lead the effort towards an ESF requirements document, and has deployed an initial tested. For Q1 2006 we expect to be deploying a first set of Edge Services for USATLAS, USCMS, and CDF within this framework on some OSG sites. The DISUN work on Condor-C in Q4 was focused on support for leases on job delegation, leading to improved support for job leases in Condor 6.7.12, released on September 28th. In addition, some design work on proxy forwarding and renewal has started.

In the area of OSG 0.4 integration DISUN has contributed by testing GIP, smrcpclient, Clarens, and JobMon. In all cases, the actual development and packaging of the tools is outside cms-dct, and we are contributing only integration effort. We deployed early versions, and validated them. Work on GIP is not yet complete, while the others are ready for widespread deployment on OSG.

1.6 Core Software and Support, CSS

USCMS has taken a leadership role in the re-engineering of the CMS software framework, following the framework review of last year and the re-organization of the CMS software effort. The status of the project as a whole is that we completed out August milestone successfully. We are now collaborating with the event filter and data management groups to satisfy their requirements. In addition we are providing consulting help for the simulation and reconstruction groups now using the software delivered in the summer. Work on the phase 2 features of the project continues.

In the software support area there has been a lot of work associated with migrating the software build tool SCRAM to the new version, and helping international CMS refine its release and distribution policies. Work to support and install the many releases needed for computing and physics challenges continue.

Software support is now providing desktop support and general software and computing consulting for the LPC. This activity has helped the users utilize the resources of the T1 group more effectively.

2.0 Core Applications and Support, CAS

Since the last quarterly report, the CAS portion of the USCMS S&C has been reorganized. Tanenbaum, Paterno and Brown now report to the Software Support area. The efforts of Tuura, Eulisse and Wildish, although tracked by CAS, are reported to the Applications Services are for Tuura and Distributed Computing Tools Area for Eulisse and Wildish. The efforts of the remaining engineers (Case, Xie, Wilkinson, Muzaffar, Osborne and Stavrianakou) will be reported here.

Xie, a developer at CERN employed by Princeton, continued her work on the POOL Filecatalog. She developed a plugin mechanism for a split catalog, and implemented a MySQL backend for the split catalog. She continued work on the RDBMS interfaces in POOL.

Wilkinson, a developer at Caltech, continued work on the Calorimetry Framework. Several new classes were added to EcalPlusHcalTowerBase, to ease its use. This class is becoming widely used by the Ecal and Hcal developers. The Calorimetry rewrite is continuing, with the intention to use as much of the CommonDet structure as possible. The Ecal and Hcal code has been decoupled, allowing one to be used without the other.

Case, a developer at UC Davis, Wilkinson and Xie are all contributing in this area. Case is working together with Muon developers to produce a design for the Muon Barrel database. Wilkinson is working together with Ecal/Hcal on their design. Both are working with the Conditions DB group led by Lueking. In addtion, Case and Xie are working together on the POOL RAL, as well as together with the LCG Conditions Database groups.

Muzaffar, a developer at CERN employed by NEU, fixed a number of bugs in the Ignominy package (dependency checking) and improved a number of algorithms. Ignominy was also upgraded to the latest version of SCRAM (the CMS build system). This was a significant effort, as SCRAMV1 has changed significantly from the earlier versions.

Case also supplied help on DDD for users and developers.

Osborne, a developer at CERN employed by NEU, and Muzaffar ran a 3 day workshop on IGUANA, to review the current applications, focusing on the most critical performance issues, 2D visualization and additional services needed. Numerous bug fixes were implemented, and many improvements made, notably in the LEGO plots and track visualization. During this time IGUANA and IGUANACMS were moved to the latest SCRAM version.

3.0 Project Office

DOE funding guidance for Tier-1, Grid and CAS efforts were stable for FY05. It was not until June that NSF guidance was firmed up. Thus for the first eight months of FY05, until the start of the new funding period, the current effort funded through the NSF stays constant, providing CAS effort and one FTE system management effort at each of the US CMS prototype Tier-2 centers, Caltech, UCSD and U. Florida.

For the DOE fraction of the S&C funds, in the following table we show the funding allocations (BCWS), the allocated effort, the reported effort and the actual invoiced costs (ACWP). We still see a large lag time between the effort being spent and the actual accounting.

FY05 Funds and Costs	FY05 total	FNAL	Universities	Caltech	Princeton	UCD	Wisc	Virginia	Iowa
Funding Allocated									
CAS Personnel	**$1,176k**	$636k	$540k		$388k	$152k			
UF Personnel	**$4,316k**	$3,664k	$652k	$376k			$102k	$74k	$100k
Tier-1 Equipment	**$1,163k**	$1,163k							
Project Office, Reserve	**$982k**	$982k							
Total Allocated	**$7,637k**	**$6,445k**	**$1,192k**	**$376k**	**$388k**	**$152k**	**$102k**	**$74k**	**$100k**
USCMSSC DOE Funds	**$7,637k**	$6,445k	$1,192k	$376k	$388k	$152k	$102k	$74k	$100k
USCMSSC NSF Funds	**$0k**		$0k						
iVDGL Funds	**$0k**		$0k						
Effort Allocated									
CAS FTE years	**9**	4	5	1	2.5	1			
UF FTE years	**31**	23	8	3			3.5	0.5	1.25
Effort Spent (12 months)									
CAS FTE years	**10.3**	5.8	4.5	1.0	2.5	1.0			
UF FTE years	**28.2**	19.9	8.3	3.0			3.5	0.5	1.25
ACWP in k$, as invoiced									
CAS Labor	**$1,368k**	$939k	$430k		$294k	$136k			
UF Labor	**$3,211k**	$2,588k	$623k	$376k			$135k	$87k	$25k
T1 Equipment	**$1,144k**	$1,144k							
Project Office	**$192k**	$192k							
Total	**$5,915k**	**$4,862k**	**$1,053k**	**$376k**	**$294k**	**$136k**	**$135k**	**$87k**	**$25k**

With the start of the new NSF funding period and the start of the DISUN project, the new Tier-2 centers start to receive funding according to a startup plan. The funding allocation of the NSF program is shown in the following table, as approved by project oversight in the PMG meeting. The Tier-2 centers receive funding through the iVDGL project, the DISUN project and the USCMS project (Research Program, RP). The table shows the allocated funds from these sources for the RP period nominally starting in May.

NSF Funds and Costs	**total**	Caltech	UFL	UCSD	UW	MIT	Purdue	UNL	NEU	UCLA
iVDGL Funding for CMS	(12 months starting 7/2005)									
Tier-2 Equipment	**$60k**	$20k	$20k	$20k						
Tier-2 Labor	**$390k**	$130k	$130k	$130k						
total	**$450k**	**$150k**	**$150k**	**$150k**						
Research Program	(9 months starting 5/2005)									
Tier-2 Equipment	**$680k**				$180k	$125k	$188k	$188k		
Tier-2 Labor	**$500k**					$125k	$188k	$188k		
Project Office										$39k
Software Labor	**$370k**								$370k	
Management Reserve	**$461k**									$461k
total	**$2,050k**	**$0k**	**$0k**	**$0k**	**$180k**	**$250k**	**$375k**	**$375k**	**$370k**	**$500k**
DISUN Project	(12 months starting 5/2005)									
Tier-2C Equipment	**$750k**	$230k	$230k	$230k	$60k					
Tier-2C Labor	**$1,173k**	$270k	$270k	$270k	$363k					
Project Office	**$77k**									$77k
total	**$2,000k**	**$500k**	**$500k**	**$500k**	**$423k**					**$77k**

The NSF funding periods cross fiscal year boundaries, while in the past we have reported within fiscal year periods. In the future we will reconcile the different funding periods in our reporting.