November 25, 2001

Draft

U.S. CMS
Software and Computing Project

Progress Report
4th Quarter FY2001

I. Summary

This document is the fourth U.S. CMS Software and Computing progress report on project status and performance to the funding agencies and to project oversight. It covers the fourth quarter of FY2001, i.e. the months July to September 2001. The project is progressing well, both in technical achievements and in management of resources.

The project was awarded an extra installment of DOE funds amounting to $285k in July 2001, allocated to equipment procurements for the User Facilities at Fermilab. Part of the labor costs for the Core Application Software efforts and User Facilities efforts at Fermilab were covered by reporting effort to the U.S. CMS Construction Project, unloading those costs from the Software and Computing Project.

At the end of September 2001, funding for a third CAS engineer was received at Northeastern University. The hiring process for this engineer is underway. The new person will strengthen the effort in the IGUANA project, WBS 2.2.

The funding situation is summarized in Tables 3 and 4 in the funding section.

II. Technical Status

In this section we report effort and technical achievements. Spent effort is assessed on the basis of regular effort reports from engineers and managers, using a bottom-up approach, where for each WBS item expenditure of labor is accumulated. For the User Facilities subproject this is reconciled with the monthly effort reports from the departments of Fermilab Computing Division. As part of the reporting for the Core Application Software subproject the engineers submit monthly reports for effort tracking.

In the following sections we use the unit of 1 FTE to denote labor of one full-time equivalent working for one year.

In the User Facilities subproject during this quarter work went on in several areas, including Monte Carlo production, hosting of the User Facility at Fermilab, support for prototype Tier-2 centers, user support, code distribution, and R&D on computer hardware. Also, project engineers helped and collaborated with the PPDG project, especially in deploying and operating the GDMP Grid Data Mirroring Package for robust file transfer and database replication between sites.

WBS 1.2 System and User Support

The Fermilab Computing Division (CD) web group continues to administer most of the project's web space, however some areas need special installations on a "private" machine. This applies to the control room logbook and the XProject servers, which are administered by the project. Also the CD-CMS web pages are continuously monitored and updated.

Work on installation and configuration of the control room logbook is ongoing. CD supports the control room logbook (CRL) and is developing a web interface to it. U.S. CMS plans to use it as a truly distributed logbook to document e.g. Monte Carlo production. Updates are being worked on by the CRL group and the project continues to follow and test new versions.

The Software Development Environment (WBS 1.2.3) includes new installations and code builds for supporting U.S. CMS code developers as well as supporting debuggers and memory leak detection tools.

Effort on the user help desk (WBS 1.2.5) is currently not billed to the project. This includes people answering the help desk phone, paging the appropriate experts and tracking problems through the system. Due to the still modest user base this activity currently is at a low level and the amount of support needed is directly proportional to the number of users.

Concerning User Training (WBS 1.2.6.1) the project is scheduling a user tutorial for later this year. Support staff took a C++ class and was trained in helping with production processing.

The prototype Tier-2 centers are now routinely supported (WBS 1.2.7) and included in the code distributions and setup of the CMS environment. Part of the support effort comes out of the User Support item in the CAS subproject (WBS 2.4).

Concerning Computer Security (WBS 1.2.8), Fermilab requires the CMS User Facilities to adhere to the security standards and to switch to a higher security level (Kerberos) from this fall on. Most of the U.S. CMS installation was converted and staff worked on security planning and implementation.

For WBS 1.2 the project spent 0.264 FTE during this quarter.

WBS 1.3 Operations and Infrastructure

Some effort was spent on software license maintenance (WBS 1.3.1) e.g. for software development tools. For Data Center Operations (WBS 1.3.2.1) the project spent in total 0.155 FTE. About 60% of a person was used for importing/exporting files from/to CERN and keeping the FNAL production status web page updated. The rest was spent on ntuple production to support the physics groups. About 1 FTE week or 0.02 FTE was spent on siting the new CMS computing equipment and getting power etc. there.

The total effort reported on this item is 0.19 FTE.

WBS 1.4 Tier-2 Regional Center

During this quarter the UCSD site of the first U.S. CMS Tier-2 prototyping center became the first regional center to export completed digitized datasets back to CERN into the user federation. UCSD also became the first Tier-2 facility to reach "Fully Operational" status with the successful test of sustainable high luminosity pile-up serving and digitizing.

UCSD has completed all assigned production requests and transferred the results to Fermilab for archiving and further processing

At UCSD a new farm configuration was prototyped in which all computational systems serve pile-up and run digitization simultaneously utilizing the Request Redirection Protocol extension for Objectivity. Tests were successful in that it solved the bottlenecks seen in the traditional configuration used by Fermilab and CERN and that it scaled with good performance on the UCSD cluster.

UCSD successfully evaluated the new version of GDMP, which can handle flat file replication. The GDMP Heart Beat monitor was also implemented to ensure that the GDMP server and clients function properly. The combination of the two has allowed the Tier-2 site to conveniently make use the Tier-1 facility at Fermilab for archiving and back-up.

Total engineering time for installation, configuration and troubleshooting/debugging was 0.125 FTE .

WBS 1.5 Networking

Under WBS 1.5.1 (Onsite Network Infrastructure) effort was spent on procuring LAN equipment for the CMS network upgrade, in specifying the equipment, following the requisition and following up with Cisco on the availability of equipment. This took about 0.5 FTE weeks during the quarter, which amounts to the total effort of 0.01 FTE.

WBS 1.6 Computing and Software R&D

Under WBS 1.6.1.2 (Hardware Testbed System) staff worked on procuring and installing the equipment for the testbed systems. Effort was spent on specification and procurement, and on hardware installation and debugging during the hardware burn-in period.

Under WBS 1.6.1.3 (Distributed Data Management) the GDMP software and the "Hear Beat" was installed and tested on the test bed system.

Some effort (0.02 FTE) was spent on WBS 1.6.2.1 (Technology Tracking and Testing).

A project, the User Analysis cluster was started under WBS 1.6.2.2.1 (Data Intensive Computing). The project is in the technology investigation phase. The MOSIX system is investigated as a solution for using a cluster of PC as an User Analysis system. This work is done in collaboration with the CDF experiment group at Fermilab. About 0.12 FTE. Has been speng on this during this quarter.

The project is investigating RAID/disk systems and is studying access pattern under the WBS 1.6.2.2.3 (Data Access and Distribution Servers). Work is going on to test disk performance and on setting up the storage hardware. We are still investigating how to best study data access patterns with the help of the Fermilab Enstore team. Detailed information on this project is available under this URL: http://computing.fnal.gov/cms/disks. The effort on this item is about 0.18 FTE.

Another work item is participation in the disk cache project. The disk cache will make data access to/from the mass storage system more smooth and transparent. The effort was spent on developing software for the cache system, kerberizing the cache ftp server, providing a Python API in addition to shell commands, and working on health monitoring tools. This system also allows the user to track data use patterns, which will be very useful in designing a data store for CMS. This is a joint project between Fermilab and Desy and is documented at: http://www-dcache.desy.de/summaryIndex.html. The project effort was 0.067 FTE during this quarter.

Additional work on upgrades to the Objectivity interface to mass storage amounted to 0.03 FTE.

Previous efforts on WBS 1.6.4 (Simulation Software) were documented, which resulted in a publication at the CHEP 01 Conference.

WBS 1.6.5 (System Administration) on the R&D cluster totaled to 0.01 FTE.

The total effort on WBS 1.6 amounted to 0.38 FTE.

WBS 1.7 CMS Detector Construction Phase Computing

For the CMS Computing Facility (WBS 1.7.1) software installation and maintenance effort is spent in the context of production, to keep code up to date by the code librarian, on installing the Iguana software, and on rewriting scripts and supporting the ntuple production on the CMS facility. This represents 0.04 FTE.

For Data Storage (WBS 1.7.2) the administrative effort for the Enstore mass storage system totaled to 0.02 FTE. Software and supporting network based file interchange (GDMP) was

deployed on the CMS facilities and the objectivity user federation was supported for distributed analysis. The mass storage interface of objectivity was supported for users with about 0.02 FTE.

On the production side of the network based file interchange work involved using GDMP to move Monte Carlo files, both FZ and Objectivity, to/from Fermilab from/to outside sites (including CERN, HIP, UCSD, CalTech, etc.). This effort was in total 0.12 FTE.

System Administration (WBS 1.7.5) amounted to 0.02 FTE.

The total effort on WBS 1.7 was 0.35 FTE.

WBS 1.8 Support for Fermilab Based Computing

There was very little effort left to cover installation of and support for desktop systems. The effort spent was 0.02 FTE.

The total FTE effort spent on the User Facilities sub-project is summarized in Table XXXXX.

WBS item	FY2001Q1	FY2001Q2	FY2001Q3	FY2001Q4	Total effort projected FY2001	Total effort FY2001
WBS 1.1 Tier 1 Regional Center					inactive	
WBS 1.2 User Support	0.26	0.32	0.20	0.26	1.31	1.05
WBS 1.3 Maintenance and Operations	0.03	0.02	0.00	0.19	0.63	0.24
WBS 1.4 Tier 2 Regional Centers	0.25	0.20	0.43	0.46	2.50	1.34
WBS 1.5 Tier 1 Network	0.02	0.02	0.01	0.01	0.50	0.06
WBS 1.6 Software and Computing R&D	0.44	0.34	0.73	0.38	2.69	1.88
WBS 1.7 Detector Construction Phase Computin	0.41	0.49	0.23	0.35	2.00	1.47
WBS 1.8 Support for Fermilab Based Computing	0.03	0.01	0.01	0.02	0.38	0.07
Total FTEs at Fermilab Regional Center	**1.19**	**1.20**	**1.17**	**1.20**	**7.50**	**4.76**
Total FTEs	**1.44**	**1.40**	**1.60**	**1.66**	**10.00**	**6.10**

Table 1 *Summary table for labor effort spent on the User Facilities subproject.*

WBS 2.1 Architecture

While the summer is traditionally one of the slowest times for progress at European based projects, there is significant progress to report in several aspects of CMS architectural development.

There is progress to report in the area of OSCAR, the CMS GEANT4 based simulation program. This year there has been an infusion of effort into the OSCAR program and it has gone from being able to do no complete detector simulation to being able to simulate a few hundred events. CMS would like to start physics validation of the GEANT4 code by the early part of next year in preparation for the physics TDR. This will require the generation of large datasets, and OSCAR will have to rapidly progress from being able to generate a few hundred events to being able to reliably generate thousands of events to stay on schedule. Hans Wenzel has been handling the OSCAR code configuration and release. Frequent release is important for fast code development.

The Detector Description Database is rapidly progressing from an interesting prototype to investigate XML, to a complete system for storing and retrieving the detector geometry information. The basic format for the XML data model is understood to store the geometry information. The geometry information will be pulled from GEANT in the short term; eventually interfaces to the detector CAD descriptions and surveyors reports may be needed. The information will be provided to CMS software clients: ORCA, OSCAR, IGUANA, and eventually FAMOS, by C++ interfaces that are currently under development. The plan is to eventually store the calibration and alignment information for the detector as well, but it is not clear what is the best method to do this. The initial prototype is expected in November. As well as development of the C++ client interfaces, there are plans to evaluate the XML storage format and calibration plans from the physicist perspective.

At the start of the quarter Chris Tully and CAS engineer Vladimir Litvin released a proposal for the redesign of the calorimetry reconstruction sub-architecture. The person who had been working in this area left CMS and development had nearly completely stopped. Several aspects of CMS simulation, including future work on the on-line development, require modification in the way the calorimetry information is calculated and stored. The redesign should also give an improvement in performance. The proposal was a phased implementation, and the release of the first phase is expected by the end of October.

The new Analysis Architecture being developed by Lassi Tuura, which relies on a very small kernel and loaded plug-ins to achieve a wide range of functionality, was presented at Computing in High Energy Physics, CHEP01. The first release is expected by the end of October.

WBS 2.2 IGUANA

Progress in the Interactive Graphics and User Analysis, IGUANA, this quarter was primarily refinement and release. Performance evaluations and bug fixes were performed. IGUANA is now at the point where is a very useful tool for debugging reconstruction algorithms by providing excellent visualization.

Work progressed on the GEANT4 visualization. There is a plan to implement an overlap finding tool developed by Martin Liendl. It will be helpful in checking the Detector Description Database and other detector geometry information by quickly indicating areas where the detector description has failed because two pieces are described as occupying the same point in space.

Ianna Osborne released a very detailed IGUANA tutorial. The tutorial gives step-by-step instructions, suitable even for very new users, that lead from installation to visualization of the detector. There was a more developer-oriented tutorial created in the spring that helped developers add objects to visualize.

There is an analysis workshop planned for the November CPT week. The plan is to have an explanation of current ideas and plans for the physics groups, to demonstrate generic analysis tools as well as preliminary integrations of analysis modules with CMS software, and to collect reaction and requirements from the physics groups to refine the program of work for 2002.

WBS 2.3 Distributed Data Management and Processing

During this quarter there was CMS planning as well as technical progress to report. International CMS' Computing and Core Software Project, CCS, consisted of six level two tasks: Computing Clusters, General CMS Computing and Software Services, Architectures Framework and Tools Kits, Software Users and Developers Environment, Software Process and Quality, and Production Processing and Data Management. CMS agreed to add a seventh level 2 task called Grid Systems, this task would focus primarily on planning and specifying requirements, evaluating and

testing, and eventually integrating Grid developed components. Hopefully this results in a more coherent and accepted long-term plan for CMS distributed production and analysis between now and the 20% data challenge scheduled for 2004. The system deployed in 2004 must be sufficiently similar to the final system, both in terms of CMS and Grid developed components, for the test to be meaningful.

There are also several areas of technical progress. The Fermilab developed production scripts have gone through several minor releases this quarter and are now documented and almost universally used in CMS distributed production. The scripts were initially developed to specify and submit production jobs for all steps of the CMS chain at Fermilab for the Fall 2000 production run. Since that time they have been increased in functionality and usability and made largely site independent allowing them to be used at many CMS production sites. This quarter the production scripts were given some automatic error recovery features and they were given a formal CMS project name, IMPALA.

There was also progress on the next generation of production tools. The project called MC_runjob is a joint development effort between D0 and CMS. It builds on the functionality of the IMPALA scripts, but it is built around a more formal architecture. MC_runjob is written in Python and relies on configuration modules called configurators. Configurators are specified for all elements of the CMS production chain and are linked together to specify a complete production job. The new system provides greater flexibility to specify production chains, a Graphic User Interface, and the ability to template production jobs. The first test of the new system will come the first quarter of FY02. The CMS Physics Reconstruction groups have requested a very large calibration sample that they would like to have simulated, reconstructed, and analyzed with only the final analysis results returned. The requested sample is 12 million events, which is estimated to take about 400 CPU months. Modifying the IMPALA production scripts to perform this task, in which most of the results are deleted, the job should run in one step, and there is a new analysis module to run at the end, would be time consuming. MC_runjob allows the configurators to be chained together to obtain expanded functionality without rewriting the pieces.

Work progressed on the new CMS monitoring and simulation. There was progress on the definition of global architecture and the implementation of the graphical user interface for configuration and control. The simulation results of MONARC simulations performed at Fermilab were presented at CHEP01. The monitoring project is very interesting because it is one of the big holes in the CMS production environment. Good monitoring tools in the short-term will allow a better understanding of the performance of the systems and allow the diagnosis of problems. In the longer-term they can improve the quality of the simulation and provide critical information to advanced grid resource schedulers.

The PPDG developed MOP distributed production prototype was able to run portions of the CMS chain this quarter. MOP was able to submit jobs from a master server to client clusters using the Globus GRAM interface. The jobs deploy the needed software using a simple software distribution technique and run locally. The results are transferred back to the master site using GDMP. The tests successfully ran CMS CMSIM tests. The more complicated CMS ORCA code, which writes into Objectivity, will be tested soon.

WBS 2.4 Support

CAS engineers have continued their efforts to support users and developers of CMS software. A few examples of support efforts from this quarter are given below. Hans Wenzel hosted a small software tutorial for users at Fermilab. The IGUANA tutorial was published and distributed by Ianna Osborne and a summer student under her supervision. Michael Case worked with the end-

cap muon groups to store parameters into an Oracle database. Vladimir Litvin provided Caltech with production support. Tony Wildish organized a Production Tools mini workshop in July, which covered a variety of development problems and plans.

The total FTE effort spent on the Core Application Software sub-project is summarized in Table 2. The total project funded effort in this quarter was 2.0 FTE.

WBS item	FY2001Q2	FY2001Q3	FY2001Q4	Total effort projected FY2001	Total project-funded effort FY2001
WBS 2.1 Software Architecture	0.44	0.38	0.56		
WBS 2.2 Interactive Graphics and User Analysis	0.06	0.13	0.19		
WBS 2.3 Distributed Data Management and Processi	0.75	0.75	0.56		
WBS 2.4 User Support	0.75	0.75	0.69		
Total FTEs	**2.00**	**2.01**	**2.00**	**8.50**	**7.75**

Table 2 *Summary table for labor effort spent on the Core Application Software subproject***.**

III. Financial Status

Table 3 shows the funding status of the U.S. CMS Software and Computing Project. Funds received in FY2000 and funds requested for FY2001 are shown. In February 2001 funding of in total $1000k was received from DOE and a loan of $500k was raised from the U.S. CMS construction project. In April 2001 the second batch of DOE funds amounting to $500k were received. A third allocation of $285k was received from DOE in July. Also $320k were received from NSF, to cover CAS engineering support, and support for a third CAS engineer of $160k was received towards the end of the fiscal year. In the table those funds are shown as “Authorized”. The budgeted funds shown are prorated for ongoing efforts.

All funds in AY$ x 1000	FY2000	FY2001		FY2001Q1	FY2001Q2	FY2001Q3	FY2001Q4
		Requested	Authorized	Budgeted	Budgeted	Budgeted	Budgeted
Core Application Software	870.0	1234.2	1234.2	268.7	288.6	288.6	288.6
User Facilities Personnel	260.7	885.4	635.4	158.9	158.9	158.9	158.9
User Facilities Hardware	344.0	1160.0	695.0	0.0	0.0	0.0	410.0
Project Office		120.4	120.4	0.0	0.0	30.0	30.0
Management Reserve		100.0	0.0	0.0	0.0	0.0	0.0
Total	**1474.7**	**3500.0**	**2685.0**	**427.5**	**447.4**	**477.4**	**887.4**

Table 3 *DRAFT! Funding status of the U.S. CMS Software and Computing Project up to FY2001Q4. All funds are given in AY$ x 1000*

There was still no news on the requested NSF funds for the Tier-2 hardware and personnel. Up to date only the NSF support for CAS engineers was received. In Table 4 the funds received and the expected funds for FY2001 are summarized.

All funds in AY$ x 1000	**FY2000**	**FY2001**		**Total**
		Requested	Received	Received
DOE	1164.7	2000.0	1785.0	2949.7
NSF	310.0	1500.0	480.0	790.0
Loan from U.S. CMS Detetector Project			500.0	500.0
Total	**1474.7**	**3500.0**	**2765.0**	**4239.7**

Table 4 *Funding sources, expected funds and funds received as of September 2001. All funds are given in AY$ x 1000*

IV. Milestone Status

1.1 tier 1 regional center

This WBS item is not active.

1.2 system and user support

No milestones projected for this quiarter.

1.3 operations and maintenance

Milestone Description	Milestone Date	Completion Date	Comments
Siting and power installation for CMS farms	10/01/2001	10/01/2001	

1.4 Tier 2 regional centers

Vendor Selection: **Expected 10/01/2000** **Completed 10/01/2000**

Final choice of vendor for all components in the first half of hardware deployment. Represents the end of performance evaluation tests as well as price and maintenance negotiations.

Procurement Phase Complete **Expected 12/01/00** **Completed 01/07/01**

Arrival of all components in first half of hardware deployment at CACR and SDSC. Racks, Computational Nodes, Data Servers/Master Nodes, Network switches, RAID arrays. This milestone was delayed due to lack of funds at UCSD.

Commissioning and Testing **Expected 02/01/01** **Completed 02/01/2001**

Of Half Size System

Systems assembled, configured, cabled, and tested. CERN standard computing environment installed and CMS software compiled and installed.

Reconstruction of events **Expected 03/21/01** **Not yet met**

from CMS Fall Production

System tested with CMSIM. GDMP used to transfer minimum bias events for use as pile-up. Production scripts deployed on tier2 site. Full contribution to CMS fall production has not been realized due to lack of manpower.

Summary of Milestones for WBS item 1.4

Milestone Description	Milestone Date	Completion Date	Comments
Vendor Selection for 1st half of T2 hardware	10/01/2001	10/01/2001	
Procurement of 1st half of T2 hardware	12/01/00	01/07/01	Milestone delayed due to lack of funds at UCSD
Commissioning and Testing of Half Size System	02/01/2001	02/01/2001	
Reconstruction of events from CMS Fall Production	03/21/2001		Milestone missed due to lack of manpower

1.5 networking

Milestone Description	Milestone Date	Completion Date	Comments
Install and commission hardware for CMS linux farms	10/01/2001	10/01/2001	
Upgrade CMS networking, Stage 1.	8/01/2001		Design complete. Waiting for equipment. New milestone data: 11/01/2001

1.6 computing and software r&d

Milestone Description	Milestone Date	Completion Date	Comments
Commissionn CMS linux farms	11/01/2001	11/01/2001	
Procurement of Test cluster	07/01/2001	07/25/2001	Bids out
Procurement of "User" cluster	07/01/2001	07/25/2001	Bids out
Procurement of additional CPU nodes	07/01/2001		Postponed
Investigating Storage (disk) Technologies	12/30/01		3Ware testing started. Other technologies ordered and installed.

1.7 construction phase computing

Summary of FY01 Milestones for WBS item 1.7

Milestone Description	Milestone Date	Completion Date	Comments

CMSIM production on central farms	12/01/2001	10/01/2001	Fermilab produced 1.4M events
Hosting User Federation at FNAL	01/01/2001	02/20/2001	Adding events as they are reconstructed.
AOD completed	01/30/2001	10/20/2001	Ntuples only

1.8 Support for host lab computing

No milestones were projected for this quarter.

WBS 2 CAS

WBS 2.1 Architecture

2.1.3.2.1.4 Release of Top Level Description Document for CMS Framework

Original Date May 1 Not yet completed

This milestone is the release of architecture documentation for the central CMS framework. The document is designed to aid users and developers who need to involved working with the CMS framework. The general lack of documentation in CMS is a problem, which will hopefully be aided by the revitalization of CAFÉ.

2.1.2.1.4 Assessment of XML Technology for use in CMS DDD

Original Date July 1 Completed August 1

This milestone is the assessment of the suitability of XML technology for use in the Detector Description Database.

2.1.2.3.5 Release of Initial Calorimetery Code Redesign

Original Date July 3 Completed July 7

This and the following milestone refer to the plan for the redesign of the calorimetry sub-architecture. There was a design document released for comments.

2.1.2.3.7 First Release of Calorimetry Code

Original Date Sep 27 Expected October 31

WBS 2.2 IGUANA

The next major IGUANA milestones are in Q1 FY02.

WBS 2.3 Distributed Data Management and Processing

2.3.7.4 Test of a Distributed Production Prototype between FNAL and U. Wisc.

Original Date July 4 Completed Oct 2

The successful demonstration of a CMS-PPDG developed distributed production system between Fermilab and a remote facility.

2.3.7.6 Test of Distributed Production Prototype between FNAL and Tier2 Facilities

Original Date August 16 Completed September 25

The successful demonstration of a CMS-PPDG developed distributed production system between Fermilab and Tier2 facilities.

2.3.6.3.11 Initial Release of Tools for Job Specification

Original Date August 6 Expected October 15

This is the initial release of MC_runjob tools for job specification and submission.