Strategic Plan for CMS and LHC (2009-2012)

Updated 12-August-2009

The CMS and LHC Strategic Plan covers the coming three years and is updated annually. The LHC strategic plans have always covered a period of development followed by an operations period with data collection. It is not the first time that CMS has hoped that the period of stable data collection is experienced before the next revision of this document, but CMS is particularly confident that this time the transition to operations will occur within the time frame addressed in the strategic plan. The Tier-1 facility has completed the extended procurement ramp arriving at the full complexity Tier-1 facility; the detector has collected cosmic data during two extended runs; the distributed computing infrastructure has been exercised at the expected scale; and while continued development effort is needed in several key areas, functional systems for early running exist.

Mission

The mission is to develop, innovate, operate, and support excellent and forefront computing solutions and services for the CMS experiment at CERN, and the U.S. CMS Software and Computing Project. In particular, we have build and will operate a Tier-1 regional computing center as part of the Worldwide LHC Computing Grid and a CMS analysis computing facility for the LHC Physics Center at Fermilab, and to collaborate closely with similar efforts at other CMS institutions in the US and internationally. It is also our goal to facilitate analysis, remote collaboration, and remote operations at Fermilab for the CMS collaboration and for the LHC accelerator.

Context and Assessment of Current State

The LHC program at FNAL is comprised of several related activities. FNAL is the U.S. host lab for the U.S. CMS project. Activities include detector maintenance and operations, detector upgrade R&D, operation of a remote operations center for remote detector operations, and data operations for CMS and for LHC accelerator monitoring. FNAL hosts the CMS Tier-1 computing facility for the Americas. The core group of application framework software developers and a significant fraction of the distributed computing and data management developers are employed by FNAL. US-CMS is a leader in the development and deployment of the Open Science Grid that benefits both US-LHC experiments and other science communities.

The LHC and CMS monitoring activities have successfully demonstrated the ability to contribute to operations at a distance. The Remote Operations Center (ROC) final facility is complete and has successfully contributed to the Magnet Test and Cosmic Challenges (MTCC); 2008 and 2009 CMS Cosmic Tests in the P5 cavern (CRAFT, CruZet); and first beam and tracker integration activities, including shift work, real-time monitoring, and data processing and reprocessing. The LHC at FNAL facility is a model for remote operation centers, and technical challenges regarding secure access to CERN-based systems are being overcome. The successful utilization of the LHC at

FNAL center has spurred the deployment of 2 other large remote centers in Hamburg and Beijing.

The Tier-1 computing facility has completed the final phase of the procurement and deployment cycle, which successfully achieved 100% scale of all important capacity metrics by the end of FY 2008. The Tier-1 facility is utilizing a node replacement model where nodes are replaced with faster hardware to accommodate the experiments need for increased capacity, while keeping the center complexity roughly constant.

Software developers at FNAL are drivers in the development and support in 3 areas critical to CMS success: the CMS Software Framework, the CMS Workflow Management, and the CMS Data Management projects.

- The CMS Software Framework was redesigned using FNAL expertise in 2006. This massive undertaking, involving coordinating contributions from global CMS, was completed on schedule. The software was successfully validated during subsequent computing and software challenges and it now in the 3rd major release. It is expected that over the time period covered in this strategic plan the framework support effort will drop significantly, but if previous experimental experience is any indicator the effort will not drop during the first 18-24 months of data collection.
- The CMS workflow management system controls the organized processing for reconstruction, skimming, and simulation at the Tier-0 at CERN, the 7 national Tier-1 centers and the 40 Tier-2 computing centers. Additionally the CMS workflow project develops the analysis tools used by the thousands of the physicists on CMS for submission to Tier-2 and Tier-3 computing centers. CMS at FNAL is not the only contributor to the workflow project, with leadership and development for analysis tools coming through INFN, but FNAL developers have acted as the architects and lead implementers for the organized processing components and several of the shared components.
- The Data Management project in CMS includes the Dataset Bookkeeping System (DBS), the Data Transfer System (PhEDEx), and the Distributed Conditions Database (Frontier). CMS at FNAL provides developers critical to the success of DBS and Frontier.

The Open Science Grid receives funding from both the DOE and the NSF for running the Grid as a distributed facility and to enable it to contribute to grid based storage, higher level distributed computing services, packaging and deployment, and support. US-CMS continues to play a leading role in the OSG, and to benefit from advances provided by OSG efforts.

Vision

The US-CMS Tier-1 computing facility at FNAL will remain the largest and most capable remote computing center for CMS. US-CMS will meet its obligations to the international experiment with processing, storage, and network resources commensurate with the size of the US fraction of the collaboration. FNAL will continue to bring to bear

the professional expertise at facility operations and deployment required to support the Tier-1 center.

By 2010, FNAL will be the best place in the world to perform CMS physics analysis. The combination of a critical mass of expertise at the LHC Physics Center (LPC), adequate computing resources to support the local analysis community, access to locally stored data samples, and a strong connection to experiment monitoring and operations will result in a world-class analysis center.

The Remote Operations Center at FNAL will provide a unique opportunity in the US for CMS collaborators to participate in detector operations and data monitoring. The operations center at FNAL will foster more inclusive participation in shift work in the US.

The software development teams at FNAL will deliver and support critical components for the experiment success in the areas of framework, workflow management, and data management. CMS will continue to benefit from the institutional expertise available at FNAL and FNAL will continue to be seen as an innovative force in HEP software development.

Stakeholders

The sponsors of the CMS and LHC work are the U.S. LHC Research Program funded by the DOE and the NSF, and the Fermilab core program.

Customers are the Fermilab and University scientists doing CMS research, and the CMS collaboration.

Effort and deliverables are provided from many groups both internal and external to the Laboratory.

Goals and Objectives

1. Operate a reliable Tier-1 computing facility that meets the requirements for custodial data storage, data serving capacity, reconstruction processing capacity, and local and wide area data serving capacity as outlined in the CMS computing technical design report (CTDR).
2. Provide a functional data operations team to peer with the CERN-based team. The CMS data operations model calls for two teams: one team operating at CERN during the European day and one operating at FNAL during the U.S. day. The teams have equivalent responsibility and authority during their respective shifts.
3. Provide sufficient local analysis computing, user storage, and data serving capacity to support a community of roughly 200 participating scientists.
4. Develop, support and operate distributed computing infrastructure to make efficiency use of the CMS dedicated and opportunistic resources through the OSG common grid infrastructure. This includes the use of processing resources for

simulation, analysis, and event reconstruction as well as the delivery of data management components to make efficient use of distributed storage resources.

5. Facilitate the development, deployment, and support of the CMS core framework to give a competitive scientific advantage to CMS in data analysis.
6. Provide opportunities for U.S. physicists to actively participate in detector commissioning and operations from the U.S. through remote monitoring.

Strategies

For the facility, CMS utilizes a strategy of active monitoring of the centers to measure readiness and reliability. Augmenting this is a detailed set of site monitoring and alarms. The CMS Tier-1 has moved from a facility doubling procurement model, to a node replacement model with the goal of maintaining a similar level of complexity. Additionally, for the last two years the Tier-1 has used a buy ahead model where approximately one third of the capacity needed in a fiscal year is procured at the end of the previous year to provide extra time for commissioning and ensure systems are in production early in the year. This will be continued as the machine schedule merits and the budget allows.

In workflow and data management as well as framework development, US-CMS has practiced active engagement with the international experiment, while maintaining a strong local core team.

CMS remains committed to the success of Open Science Grid and maintains close integration with OSG development and deployment projects.

Resource Needs

We expect the engineering resource needs profile to slowly decrease by 2-3 FTE total through these three years, in terms of the overall FTE count, due to flat funding. There is the possibility of compensating effort through the CMS upgrade project. The biggest change will be the focus of the supported effort from roughly equal amounts of operations and development effort currently to a more solid majority in operations. In the first 2 years of operations we expect the development needs will stay high, with more significant reductions in year 3. The physics effort will increase and additional scientists, RAs and guest scientists will be required to start the physics program of the LPC.

The funding profile for equipment procurements will average to about flat with some structure year-to-year based on the machine schedule.

Progress Indicators

CMS will transition from a preparing to a running experiment beginning with the start of LHC operations in November 2009. The goal for CMS is to achieve a level of service reliability and transparency in the distributed services comparable to what can be achieved with local access to computing resources.

By January 2010

- CMS expects be able to successfully execute 25,000 jobs per day on the FNAL batch computing resources through a combination of local submissions to the LPC and grid submissions to the Tier-1.
- The local community should be able to submit 20,000 jobs per day across LPC analysis and remote resources.
- CMS expects to be able to reliably accept 400MB/s of data from CERN to tape.
    - Time from deciding to transfer data to achieving the desired throughput should be less than 1 hour.
- The goal is to provide transfers from FNAL storage to Tier-2 centers at between 50MB/s, to the worst connected Tier-2 centers, and 500MB/s for the best connected Tier-2 centers.

By July of 2010

- CMS expects that the number of job submissions to FNAL and to U.S. distributed resources will increase slowly, but that the volume of data being accessed and the diversity of samples requested will increase.
- The Tier-1 center will execute at least 4 reprocessing passes of the custodial raw data.
- In 2010 the accelerator is likely to be collecting low luminosity data until April, but ramping toward higher luminosity. The physics groups at the LPC will have access to sufficiently large samples that analyses of rarer signals will be possible.