Run II Computing Review

September 11-12, 2003

Executive Summary

The second review of Fermilab Run II Computing took place at Fermilab on 11-12 September, 2003. The members of the review committee were:

- Ian Bird, CERN (chair)
- Elizabeth Buckley-Geer, FNAL
- Bruce Gibbard, BNL
- John Hobbs, Stony Brook
- Stephen Kent, FNAL
- Dave Morrison, BNL
- Christoph Paus, MIT

The charge to the committee included the following: consider the mode of operations and suggest possible improvements through consolidation of efforts; assess whether the experiments understand their resource requirements for operations and future evolution; consider the projected costs and the cost model; address whether the data handling systems are appropriate and aligned with grid computing developments; and assess whether the management and communications between groups and departments involved are appropriate and working.

The review team heard presentations from members of both CDF and D0 collaborations as well as representatives from the Computing Division (CD).

It was clear that both experiments have made tremendous progress since the first review and are actively and successfully recording and analyzing data and presenting results at conferences. The committee congratulates the experiments and the Computing Division on their success, and it was clear that many collaborative efforts had played a significant role in this.

The data handling systems from end to end are working well and in particular the Enstore/dCache system shows a high level of reliability and performance. The move to the use of SAM by both experiments was beneficial and CDF is encouraged to pursue this path more aggressively. This might be aided by making SAM a division-level project.

The cost drivers for data handling are very different between the experiments. For D0 the reconstruction time is the limiting factor, while for CDF the proposed increased data rate drives the cost. For D0 it is important to significantly reduce the reconstruction time. However, the proposed budget from CDF would make its request significantly greater than that of D0 even with D0's poor reconstruction performance.. It was felt that in the light of the large requirements produced by CDF's increased data rate that CDF should be more proactive in finding external computing resources to address part of the cost.

The operations model seems to be quite effective. However the committee felt that there was potential for improved operations efficiency in several areas. These included consolidation of management of the farms, of consolidating Enstore instances, and potentially further consolidation in the database area.

Summary of Recommendations

- CDF should more aggressively pursue a full SAM implementation plan

- CD should consider making SAM more of a division-level project, and aim for an appropriate balance of effort from the two experiments
- CD should consolidate the administration of the Enstore systems
- CD should work on a risk assessment and benefit analysis for the proposed data dispersal plan
- CDF should invest effort in a “thumbnail” data format
- D0 should further review its reconstruction code including profiling and comparison with Monte Carlo and L3 code
- D0 should resolve the system administration problem for the clueD0 cluster
- We recommend aggressive removal of the D0 SGI system
- CDF should clarify the projections for their computing needs which depend on a model that underestimates true usage by factor of 2.
- CDF should take into account the international contributions to CAF and remote computing.
- CDF should find remote facilities to help meet the demands of the proposed data rate increase
- The experiments should recognise that database technology is now a key technology, and consequently should ensure that they provide sufficient support
- There should be more coordination and coherence of database related developments
- CD and the experiments should seriously investigate forming a joint operations group for system administration of the farms.

Data Handling

At the time of the Run II computing review conducted in 2002, it seemed that the Fermilab data handling systems would be capable of handling the anticipated needs of both CDF and D0. A bit more than a year later, it is clear that the data handling systems have been demonstrated very successfully to be capable of meeting the needs of the experiments. The committee congratulates the teams involved on all sides (experiments and CD) on this significant achievement.

The CDF plan for a substantially increased data rate is the most significant driver behind the projected need for a more capable data handling system in the near future. It is noted that, although impressive steps have been taken to test and stress the capability and scalability of the system deployed today, it cannot be assumed that the system will scale to the needed size without the investment of substantial manpower. This requires a more detailed understanding of what resources will be required to achieve this.

The Enstore mass storage system has proven itself capable and stable, with just over 1 PB currently archived on tape and moving between 10 and 20 TB of data per day. There is a clear plan to migrate from the older STK 9940A drives to the newer STK 9940B variety. The facility expects to use the less expensive LTO2 drives and media in addition to those based on STK 9940 technology. It's not clear that one would recommend from the start acquiring multiple tape silos in order to use two different tape technologies; however, given that the different existing silos are already configured to use these different types of tape drives, it seems a sound strategy to continue with the two varieties for the time being, balancing the known capabilities of the STK drives while gaining further experience with the LTO2 drives.

Other plans for the mass storage system were outlined; in particular, a plan for the consolidation of the administration of the three extant systems and a plan for catastrophe-mitigating data dispersal. The benefits of the former seem clear and the committee fully endorse proceeding with this plan. However, the committee felt that a clear case needs to be made for the benefits of a data dispersal strategy with a more thorough risk and cost assessment before launching a major effort in this direction.

The dCache disk caching software has also proven itself quite a success during the last year, though it has taken very significant effort within CDF and CD to undertake the development necessary to turn dCache into a production-ready system. In addition to easing the management of large quantities of disk, it is hoped that with the deployment of sufficient disk, a "tape less" analysis path will become possible.

The plans by D0 to implement "thumbnails" of approximately 10 kB/event has been largely realized, with analyses based on these data structures being the usual approach pursued in the collaboration. The committee noted that CDF do not have an equivalent to this thumbnail dataset, and the committee felt that CDF should be encouraged to invest some effort into producing one as there seem to be clear benefits possible.

SAM, the system for "Sequential Access via Metadata", initially developed within D0 has continued to serve that experiment well, and following the investment of significant effort by some of the key developers to adapt the metadata schema it also promises to do the same for CDF. However, the state of SAM adoption by CDF seems to be little changed since the 2002 Run II Computing Review. The development of SAM seems to have been slowed because of the diversion of effort to reconcile the CDF and D0 schema, so it would be advisable for CDF to move more aggressively to adopt SAM. This would help solidify the status of SAM as a division-level project and aid CD in aiming for an appropriate balance of development and maintenance effort from CDF and D0.

Farms and Production Facilities

Farms are now not only used for data reconstruction but are playing an increasingly major role in analysis computing. Both CDF and D0 now have substantial analysis farms in addition to their reconstruction farms. We encourage both experiments and CD to explore ways of consolidating the management of the various farms as much as possible. They should look at the technical issues involved and propose a schedule for implementation.

Reconstruction farms

Both experiments presented projections for their reconstruction farm needs through FY06. The two experiments have different cost drivers. For CDF the proposed rate increase due to the CSL upgrade keeps the costs constant even though the reprocessing fraction drops to 0.3 in FY06 due to increased stability of the code over time. For D0 the speed of reconstruction is the major cost driver. In both cases the cost model looks reasonable, assuming the inputs remain unchanged.

The biggest concern is the speed of the D0 reconstruction code. At the review in 2002 the current execution time was 10 GHz-sec/event including ntuple creation. A worst-case scenario of 25 GHz-sec/event was used for planning purposes at high luminosity. At this review we were shown times of 50 GHz-sec/event at 4×10^{31} cm^2 s^{-1} with a clear tendency to increase with instantaneous luminosity. There has been an important increase in the efficiency of the fiber tracker algorithm from 75% to 98% which is good. However the current reconstruction time now exceeds the previous worst-case scenario by a factor of two. If this number continues to increase then there may be consequences for the D0 physics program as the ability to reconstruct the data within available budgets diminishes. We note that the execution time drives the need to rely heavily on remote facilities to provide the reprocessing capacity. We were shown that a reduction of 10-20% in execution time could be achieved by skipping events that have very high occupancy in the fiber tracker and take a very long time (~10^5 seconds). This is also contributing to a reduction in the efficiency on the farm. Ideally one would like to be able to achieve more like a factor of 5. Obviously there is some limiting factor in speed beyond which one cannot improve due to the nature of the tracking detector but we feel that we are quite far from that limit at present.

We strongly encourage D0 to do a full code review including but not limited to profiling, comparisons to Monte Carlo data and the code running in Level 3 (which it was noted by one of the reviewers runs much faster).

The largest cost driver for CDF is clearly the proposed increase in data rate. From what the committee could learn, given such a proposed increase, the proposed cost model used by CDF is reasonable, as are the assumptions of the reprocessing model.

Monte Carlo Farms

D0 has continued to produce all of their Monte Carlo events at remote facilities. In the last 2 years they have produced about 80 Million events. This is an effective use of off-site resources and they are to be commended for this.

At the last review CDF had no plans to generate Monte Carlo at remote facilities. At this review we were told that there are now substantial amounts of Monte Carlo being generated off-site. A total of 110 Millions events were generated in Canada (60%) and the UK (30%) and there is an MOU with Canada for production coordination. They are also working on making a general off-site user facility for Monte Carlo generation. This is a very positive development and CDF is to be congratulated on having taken this step.

Central Analysis Facilities

The past year of operation was the first with significant collider data taking. Many aspects of the plans presented to the last review were successfully proven. Both the CDF Central Analysis Facility(CAF) and the D0 Central Analysis Backend(CAB) have been quite successful. Linux-based file servers planned by CDF have been implemented, and they are now performing as expected after significant effort to understand and resolve hardware and firmware incompatibilities and problems.

The committee also notes that the D0 desktop cluster maintained by the experiment, clueD0, has been a striking success, providing a batch system and data access via SAM. This enabled the efficient use of the large amount of compute power on individual desktops for analysis and simulation. At the time of this review, the future management of clueD0 is unclear. We encourage the experiment to resolve this as quickly as possible.

The majority of CDF CPU power and budget costs are used to provide the CAF. This seems reasonable given the experience of the past year. However, projections for future needs are based on a model for which the FY03 CPU prediction fell short of the actual need by a factor of two. The experiment feels this was caused by anomalous use of CAF for Monte Carlo generation but has not demonstrated this clearly. Because the CAF drives the CDF computing model, the committee strongly urges that CDF confirm the validity of the prediction.

The initial D0 CAB system has been brought to fruition during the past year. Given the success of CAB and of the CDF file server model, the committee recommends an aggressive schedule for the decommissioning of the SGI SMP system, d0mino, moving the event picking and disk serving to Linux systems. There is no longer any good technical reason to maintain the SGI system longer than necessary.

Remote Analysis and Grid Computing

Both CDF and D0 experiments have acknowledged the importance of using remote computing resources. The usage of the remote resources ranges from the simplest applications like Monte Carlo event production to the most complex ones like data re-processing and user analyses. Since the last review significant progress in implementing the usage of offsite computing has been made.

D0 had already produced large Monte Carlo samples offsite at the last review and has continued this effort. The experiment has now started to re-process data offsite as well. The full integration of SAM into the D0 computing framework has been essential for this effort.

At the time of the last review the usage of remote resources by CDF were rather limited. Since then CDF has produced all official Monte Carlo at remote facilities which is a major step. The full integration of SAM and the related tools into the CDF effort is essential to exploit the remote facilities reliably. As the data handling situation with dCache and Enstore has stabilized the committee recommends the CDF team to follow up on SAM integration aggressively. The committee congratulates both experiments for their progress and encourages both experiments to expedite implementation of their further plans.

For CDF the additional computing needs due to the larger data volume is a good candidate for pushing the remote facilities. In the accounting of the computing resources two related issues have been noted by the committee. The CAF CPU procurements from non-Fermilab sources were significantly reduced for the years 2004, 2005 and 2006. At the same time no additional remote resources have been accounted for. This strikes the committee as conservative. We encourage CDF to review the schedule for including remote facilities particularly in view of the proposed additional data volume from the increased trigger bandwidth. It seems feasible to the committee that remote computing power could and should be included with an appropriate efficiency factor in the accounting.

Due to the long event processing time, D0 depends critically on remote institutions for the data reprocessing. It seems essential to validate this procedure for large scale production by the fall. From the documentation provided it is not obvious that the remote computing capacities are sufficient to accommodate the D0 needs.

Since both experiments will depend to a certain extent on their remote facilities, the committee re-iterates the importance of formalizing the process of including remote resources into the computing plan. Therefore the requirements for remote facilities to be counted in the respective computing plans have to be exactly specified and agreed upon in Memoranda of Understanding (MOUs). In particular the type of usage of the facility, its maintenance plan and the amount of CPU and disk space should be included.

The grid-enabled analysis tools in CDF and D0 are intimately connected with the SAM integration and a lot of progress in both experiments has been made. The development of SAM into a grid-enabled tool is being pursued by a group of people from D0, CDF and the computing division and has been commented on before. It is quite obvious that the grid is an excellent tool to smoothly and effectively include resources from remote facilities at all levels and should therefore be strongly supported by both collaborations.

Infrastructure

Networks

The usage of remote facilities will inevitably require strong wide area network (WAN) connectivity. The committee recognizes the positive development which has been made in the last year. The ESNET link has quadrupled, as planned, and the connection to STARLIGHT has been established on a general level. STARLIGHT provides access through Internet 2 and CASNET which is available to the North American and European institutions. This development should be continued and STARLIGHT should become a fully integrated part of the WAN plan to allow for expedient shipment of data between the different sites. The committee endorses the plan to pursue further upgrades to the ESNET link.

The plans presented for network infrastructure upgrades within the lab seem reasonable and in line with accepted standards. In particular, providing Gigabit links to certain desktop systems (providing some justification of need) is now reasonable, with wider deployment coming later as part of the general infrastructure upgrade planning.

Databases

The committee noted that good progress has been made in making DB systems more scalable and robust. However, there seems to be an insufficient level of effort provided from the experiments to support the database systems. The committee recommend that the experiments must recognise that database technology is now a key technology, and that they rely heavily upon it, and that consequently they should invest effort to ensure that they provide sufficient support. This is clearly also an area where there is much opportunity for coordination and coherence of efforts.

The committee felt that the current strategy of using Oracle for production database services, and developing the use of OpenSource systems like Postgres or Mysql for remote sites seems reasonable. On the issue of the Oracle licensing costs, the committee agree that the current philosophy of closely monitoring the licensing costs is appropriate, and that should the costs rise significantly the true costs of developing an alternative service based on another database system must be fully understood before switching. Such costs will be significant in developer FTE and essential operational developments.

Staffing and Budgets

Staffing

The committee found it difficult to glean from the various presentations an overall understanding of the way in which effort is being allocated. For example, in light of the fact that Run II is classified as the Laboratories highest priority activity, it was surprising that only 11 out of the 56 people (<20%) in the Computing & Communications Fabrics Dept. were chargeable to Run II experiment support. In addition, no detailed breakdown regarding the assignment of personnel within the CDF Dept. was presented. While something less than a full-blown resource loaded WBS for the Run II computing project seems appropriate, a table which allows one to immediately see globally how much effort is being devote to which activities for each of the experiments would allow one to make more informed judgments regarding the relative cost/benefit of various personnel allocations.

It was clear from the presentation made that personnel in many areas were stretched thin and that coverage in certain specialties was only 1 person deep. However, it appeared to the committee that there was the potential for personnel savings by combining certain areas of activity; among the most obviously being the combination of the various farm administration activities into a single effort and the retirement of legacy SGI support.

The committee felt that since SAM is now being used by both CDF and D0, its development and support should be moved from the D0 Dept into a Dept. within CD which is non-experiment specific as is currently the case for Enstore and dCache. This would establish SAM as a truly shared product and clarify attribution of its development and support costs.

The committee was impressed by the modest level of development and support, ~9 FTE's, required to deliver the high levels of functionality and reliability, which have been achieved with the Enstore and dCache systems.

Budget

The capital equipment budget for D0 is dominated by the CPU cost of re-reconstructing data primarily because the reconstruction program is very slow. As mentioned elsewhere, a serious effort should be made to improve reconstruction performance or to prove that it is in fact a limitation of the detector. Failing this, it seems appropriate, as D0 has done, to relate the level of re-reconstruction (frequency) with the amount of additional computing capacity that can be delivered by outside collaborating institutions.

The capital equipment costs for CDF are driven by the increase in disk and analysis CPU costs associated with the new plan to substantially increase the data taking rate. It was not immediately obvious to the committee why this increased data taking rate did not significantly increase the reconstruction CPU cost. In the context of this change in planned data rate, the committee felt it would be appropriate to open the issue of contributions of compute capacity from outside collaborating institutions.

While there was considerable discussion of budget from the perspective of capital equipment, as was only briefly mentioned in the for CD presentation, the total cost associated with computer support of the Run II experiments is dominated by personnel cost, $7-8M for personnel versus ~$3M for capital equipment. In trying to evaluate the appropriateness of the budget plans and the trade offs which are being made, it would have been very useful to see a more comprehensive costing including personnel, service and maintenance contracts, supplies, etc. Perhaps for the next review such a costing could be prepared. However, it is clear that if there is a significant budget shortfall for Run II computing, a global optimization including all cost terms should be done. In such an optimization the leverage of a factor of 2-3 times higher personnel cost may be important.