ATLAS Computing Overview Day **CERN, December 3, 2003**

COMPUTING PANEL REPORT

J.J. Blaising, L. Fayard, J. Huth (Chair), M. Kasemann, I. Videau[1]

Introduction

The computing panel was established upon the request of the Atlas Deputy Spokesperson T. Åkesson. The main charge to the panel was to take stock of where ATLAS computing stands, and identify interface between detector subsystems, trigger and data acquisition and the LHC Computing Grid Project. March 2002 was the last review.

There were 22 presentations in all. Most of the presentations were quite clear, and covered the following areas:

Management
T-DAQ
Core Software
Reconstruction
Simulation
Computing model and grids
Calibration and alignment
LCG Applications
Data Challenges
Combined test beam

Each of the speakers was asked to address the following points:

What is the scope of the task?
What are the goals?
What is the situation with human resources?
What has been achieved?
What kind of user feedback have you received?
What are the main interface areas in ATLAS and LCG?
What are your main concerns?
Where are there areas for ATLAS coordination?
Where are there areas for LHC-wide coordination?

Most speakers addressed these questions, although some deviated from this format.

[1] Torsten Åkesson, Deputy Spokesperson for ATLAS, served ex-officio

The overall impression was that a large amount of progress had been made since the last review in many areas. Some areas that had been cited in the last review were still sources of concern, and with the progress made in many areas, concerns about new areas, particularly in the interfaces with database systems arose.

The main concerns the committee had were:

Human Resources : The shortfall in manpower on ATLAS will require continued effort to get human resources while making appropriate adjustments to schedule and scope, priorities and overall coordination to achieve a good fit between the goals and resources available. Software Infrastructure should have the highest priority for added FTE's. In particular, the combined test beam has a fixed schedule and should be considered at higher priority than DC2, although elements of the computing model should still be pursued well in advance of DC2.

Database Coordination: There is a need to have more widespread coordination across the areas of the database: event store, conditions, calibration, and technical coordination. Although there is a coordinator and a task force with a charge to look at this coordination, it is important to draw more people into this area.

Detector Geometry:The input to the GeoModel for detector description comes from a number of sources and needs to be established in a standard, maintainable format.

Language Support: The committee is concerned about the amount of effort it will take to support languages other than C++.

Reconstruction Coordination: The committee is concerned that there needs to be a sustained effort to coordinate reconstruction across many areas. The concept of a "rotating chair" will not be effective, in our opinion.

Graphics: The committee was concerned that there was a multiplicity of graphics options. Moreover, there didn't appear to be any requirements existing for the scope of graphics support.

Below, we have discussion on these and other points, and call out recommendations with bullets.

Management

Human Resources

The human resources situation for ATLAS software was reviewed by the LHCC in September 2003. At that time ATLAS reported a critical shortfall of software developers and physicists involved in software development of about 35%. First steps to solve this

problem were taken, but the situation has not changed much and is a serious point of concern. The panel makes the following observations and recommendations:

• Rec1: We support the plan to establish a core team at CERN to increase efficiency of development and integration work. SIT should be the highest priority of new hires. ATLAS should be proactive in approaching funding agencies in different countries to try to fill remaining gaps.

• Rec2: We recommend that ATLAS should strive for a closer coordination of infrastructure and software support between TDAQ, HLT and the corresponding offline activities to maximize synergies in these areas.

We commend the planning activity together with regular status reports from all activity areas. It has shown to be a valuable planning and monitoring tool. We encourage the integration of the offline planning into the PPT planning tool and the definition of dependencies between other ATLAS areas, especially installation and commissioning. The schedules and milestones should be regularly updated and kept realistic to reflect the available human resources.

• Rec3: In particular the currently planned schedules for the CTB and DC2 seems to be unrealistic, as a side effect it may seriously overstress the key people involved. It should be re-considered and updated to reflect the expected readiness of the components.

• Rec4: We consider the successful support of the CTB of highest priority. It is an important milestone for the subsystems and it is expected to bring vital feedback for the offline software deployed from a broader range of users.

Platform support

The panel observes, that ATLAS (like several other experiments) currently develops and supports their software mainly on one combination of platform, operating system and compiler. Although this reduces the effort on the short-term, the panel believes, that such a situation has significant draw-backs on the long-term in the areas of code robustness and lack of deployment and exploitation of computing resources on the Grid.

• Rec5:We recommend that ATLAS develops a realistic policy for the evaluation and the support of new platforms and increases the number of supported platforms to maximize the benefits listed above.

Language

The panel heard about the continued use of Fortran90 code in the Muon-subsystem and the proposal to develop new Fortran90 code for the AMDB. We have serious doubts that Fortran90 is a viable option in view of long-term (>10 years) availability and cost of support. Our main concern is the availability of the Fortran90 environment and the

support load for a Fortran90 environment at all levels: development, interfacing, maintenance, code distribution, training, documentation, deployment on the Grid and offline code re-use in ATLAS online. Beyond this, the coupling of other applications, such as the HLT to Fortran90 may create unsupportable dependencies.

• Rec6: ATLAS should re-examine the support of non-C++ code and the issues of transportability and deployment of other languages in light of further experience in code installation in remote sites and an understanding of HLT requirements.

DC2-TDR-MOU Scheduling

Significant milestones in the near future are the definition of the ATLAS computing model and the computing TDR together with the MoU for computing for ATLAS and the LCG. The proposed schedule seemed inconsistent to the panel, although it did not understood in detail the dependencies and sequence of these important events. We recommend that ATLAS makes an attempt to clarify the sequence and schedules together with CERN management and LCG.

• Rec7: ATLAS should rationalize the schedule of DC2, computing model, TDR and MOU preparation.

A task force has been created that will advise the CERN management on the structure of memoranda of understanding between CERN and funding agencies and participating institutes. This is expected to report on a timescale of roughly October 2004. It is foreseen that there will be five types of MOUs. Four of these will be for computing needs of each LHC experiment, and one will pertain to LCG activities.

•Rec8: The committee believes that the memoranda of understanding should be phrased in terms of deliverables and services, rather than a raw accounting of human resources of cash-equivalents. From this standpoint, the computing MOUs would have a similar structure to the detector construction MOUs.

Computing Model, grids, Data Challenges

Distributed computing in a grid context is particularly complex due to the fact that it requires the integration of software developed by different providers: ATLAS, LCG, etc, with interfaces to different kinds of middleware. It must be installed and maintained over many different sites in a heterogeneous environment. Databases must be installed and kept current. Many users performing analysis must be supported and resulting data published to the entire collaboration.

The user requirements have been captured in the HEPCAL I and HEPCAL II documents. Input from physicists is critical to the proper design of high-level application interfaces. It is also highly important that ATLAS provide direct and coherent feedback to

middleware providers. The issue of database scaling is a particular concern in this context.

Data Challenges

Data challenges are essential for assessing the computing model. Experience gained from DC1 has been factored into some of the requirements for DC2. A new element has been introduced, however, into DC2: that of distributed analysis.

- Rec9: Develop a plan to evaluate and compare the performance of grid tools during DC2

- Rec10: Develop a strategy for grid monitoring and reporting.

- Rec11: Understand interoperability of grids at a deeper level than what has been described in the presentations. In particular, we are concerned that ATLAS needs to examine how different grids function at a deeper level than putting a simple common interface layer on top.

- Rec12: Develop a plan for DC2, Phase III for testing analysis activities. In particular early tests of distributed analysis are highly encouraged.

- Rec13: Work with grid developers to establish mechanisms for better stability and robustness of grid operations in a production mode.

Distributed analysis

Although stable production running in grids is only possible with large amounts of human intervention, it is not too early to plan for the needs of distributed analysis. The committee was happy to hear that ATLAS had created the position of a distributed analysis coordinator. The ARDA project launched by LCG is supposed to deliver a prototype for distributed analysis on a timescale that may be beyond DC2, but nonetheless is a potentially interesting direction with the adoption of OGSA standards for grid middleware.

- Rec14: It is important that ATLAS be fully involved in ARDA development process to ensure that it reflects the needs of the collaboration.

Computing model

Certain aspects of the computing model have been around for some time: the CPU requirements, storage requirements etc. These have been progressively refined and updated. The computing model calls for a "cloud" of Tier 1 centers that share resources. A fraction of the ESD data ($1/3^{rd}$ approximately) will be stored on disk at each Tier 1 center. Other issues, such as distribution and installation of code, and the

distribution of configuration and calibration data in a timely way have not yet been addressed.

• Rec15: The computing model should be updated to reflect the needs and manner of distribution of data, particularly configuration and calibration data.

Coordination of online/offline activities

Common tools and practices for offline/online/TDAQ

The panel acknowledges the joint effort done recently by the offline and online groups during the writing of the HLT/TDAQ TDR . However he is worried by the future relations , in particular because several physicists have left the TDAQ effort after the completion of the HLT/TDAQ TDR . We encourage a restructuring of the reconstruction effort that could help common work but we are worried that nobody is willing to handle this effort

• Rec16: Continue to work to identify an overall coordinator for reconstruction over all areas.

Database efforts

Much work has gone into the definition of the event store database. Other areas, such as configuration, TC, and calibration database are being used with a mixture of solutions. The configurations and calibration database are very important for the Combined Test Beam. An overall coordinator has been appointed, and a task force formed to look into common needs in the database effort. We applaud this move, but note that the effort level is going to increase tremendously as the turn-on of the LHC approaches.

• Rec 17: The panel recommends the development for a plan for rationalization of all databases for TDAQ and offline and suggests a stronger organization of this effort with sufficient manpower.

The information that is used for the geometric description of the detector is presently created and stored from a number of sources. This includes G3 volumes, C++ code, text files, F90 code and a number of other sources. Ultimately this is stored in the Nova database, and then converted into the GeoModel. The multiplicity of sources of detector geometry information is a concern to the panel. In particular, the transparency of access to information and the long-term support of this information is a major concern. While the panel does not wish to impose a standardization that extends deep into the construction details of detectors, we feel that it is important that a uniform expression of geometry information be established and a uniform way of populating the configuration database be made.

• Rec18: Create a uniform definition of geometry information, and supportable method of loading and accessing the configuration database.

Core Software

A significant amount of progress has been made in core software since the last review. In particular the incorporation of Geant-4 into Athena, and use of Athena for physics studies and the HLT TDR is very encouraging. The GeoModel appears to be a well-supported and viable means to describe the detector in the context of Athena. The POOL database has been incorporated into Athena and it now important for the collaboration to gain some experience with it.

Event Data Model

There is a concern about the speed of progress of the Event Data Model (EDM). This is a crucial item, in particular for the combined test beam. As a result of deliberations by the reconstruction task force, there was an overhaul of substantial portions of the EDM. While the structure is now better defined, there is an immediate need for a definition of priorities for 8.0.0, particularly in the inner detector area of track class definitions.

• Rec19: The software management team should spell out the priorities and essential milestones for implementing the RTF recommendations for release 8.0.0.

ESD and AOD definition

The ESD and AOD formats have not been fully defined as of yet. Although one might be able to guess at the rough structure of these, the question of navigation back to raw (and truth) data is important since it gets tied back into issues of the computing model and data challenges.

• Rec20: Identify critical aspects of the ESD and AOD definitions that require coordination with the computing model.

Graphics

Graphics, and event visualization in particular are essential tools for debugging algorithms and analysis efforts. A large number of graphics packages were presented to the committee, although it appears that the default option to be supported is Atlantis. The committee was pleased to see that new effort is likely to come from the UK for graphics, but is concerned about the multiplicity of packages and the need for a clear set of requirements.

• Rec21: The committee recommends that only one graphics package be supported and that a set of requirements be established in coordination with the user community.

User Feedback

A number of comments were made during presentations about a deficit of user-feedback to developers, despite many attempts to solicit input. The joint reconstruction effort and combined test-beam should help provide more feedback to the developers, but there should also be an active engagement of the physics community in providing feedback.

- Rec22: The physics community needs to engage itself in providing feedback to the developers.

Reconstruction and Simulation

Much of the above discussion has items pertaining to reconstruction and simulation and won't be repeated. The committee was pleased to see progress in a number of areas: Geant4 and increased usage of Athena.

Combined reconstruction

As the software matures, and with the approach of the combined test beam, there is a need to address reconstruction issues across detector areas. Beyond this, the definition of physics objects begins to get tied into this area. The computing management had tried to establish the position of an overall reconstruction coordinator, but could not find anyone suitable who was willing to take on the job. The present arrangement is to have a rotation of chairs of the joint reconstruction group. The committee is concerned that this may lead to a lack of organization and progress in this important area.

Distributed Analysis

With the progress in production-level software, more attention is now being focused on the support of data analysis. In the context of grid computing, this now becomes "distributed analysis". Already there have been efforts to create a framework to support distributed analysis. In particular, for ATLAS, there has been the creation of GANGA and DIAL. Recently, an LCG RTAG has recommended a project, ARDA, directed toward distributed analysis. Although the original mandate of ARDA was to establish the needs of distributed analysis, the project created in the wake of ARDA was to employ OGSA standards to reengineer the AliEn package to produce a prototype distributed analysis framework.

The committee is pleased to see that the computing management has established a person to oversee distributed analysis for ATLAS. It is important that this person get involved in the ARDA process to ensure that the analysis needs of ATLAS are properly reflected.