Peter Loch
University of Arizona
Tucson, Arizona
September 8, 2009
Draft 1.0

Summary of the US ATLAS Software Planning Retreat

University of Arizona
Tucson, Arizona

August 10-11, 2009

1. Attendants

Paolo Calafiura[1], Jim Cochran[2], Kaushik De[3], Peter Loch[4], Fred Luehring[5], David Malon[6], Eric Torrence[7], Alex Undrus[8]

2. Preliminaries

2.1. Content of this note

This note does not actually provide exact minutes of the retreat, rather it is an attempt to summarize the most important outcome of the discussions together with some updates and afterthoughts mostly on details of implementations. Some aspects of discussions after the meeting, mostly with Srini Rajagopalan and Jim Shank, are also included.

2.2. Status

Draft 1.0: *first draft by PL – no feedback from attendants included.*

3. Meeting goals

The main goal of the retreat was (1) to discuss priorities in software and (2) develop an improved management plan for the US ATLAS software effort (WBS 2.2). In particular, the implementation of the recommendation from the May 2009 US ATLAS software review stating

[1] LBNL
[2] Iowa State University
[3] University of Texas at Arlington
[4] University of Arizona
[5] Indiana University
[6] ANL
[7] University of Oregon
[8] BNL

In general crisp, clear, prioritized lists of activities and tasks will facilitate communication of Software's efforts to both global and US ATLAS management,

is to be discussed. More specific goals for the software discussion (1) are

- The review of activities and scopes in the context of the next fiscal year (FY10);
- Establish prioritized lists for short and long term efforts according to recommendations;

For the management process (2), the items to address are:

- Re-structuring of the WBS;
- Reporting at bi-weekly US ATLAS/US CMS joint agency phone calls;
- Resource loaded effort tracking and corresponding streamlining of the quarterly report structure;
- Decision process on scope changes and/or (short and long term) extensions and the corresponding re-direction scenarios;
- Organization of milestones and general reporting system improvements;

Not all of these goals have been completely or sufficiently addressed at the retreat. This brief report focuses on the ones which have been discussed in some detail. Additional points of discussion include setting up resource loaded effort tracking for each FTE – these have been addressed (see below).

4. Management issues

The discussion of the software management plan preceded the software discussion. The issues discussed in detail and the conclusions are given below.

4.1. New WBS structure

Following the suggestion by Srini Rajagopalan we discussed a new WBS level 3 and 4 structure for WBS 2. The result of this discussion, including the now incorrect assumption that all software efforts related to distributed computing and analysis support are moved to their own respective level 2 category[9], are shown below. In any case, we suggest adding a "Maintenance, Operations & Testing" level 4 category to each level 3 category. Nearly all supported FTE spend a considerable amount of time on this issue, and this effort needs to be reported as well. Alternatively, this report could be attached to the level 3 subject categories reports, but still needs to be identified as such[10].

Action item: *PL to discuss new WBS with Srini and Jim, with the implementation of final agreed-upon version before end of September 2009.*

4.1.1. WBS 2.2.2 – Core Software

[9] The recent draft of the PMP indicates that the distributed computing effort is merged with facilities (WBS 2.1). This is of no consequence for the re-organization of WBS 2.2 discussed here.

[10] Any software deliverable is considered to bring along an M&O responsibility after development and implementation. Assuming the responsibility to produce the software implies this additional effort even beyond the actual delivery date. We should know about the amount of effort and the time spent on it, especially versus time spent on new developments.

The table below shows the new suggested new WBS structure for core software, without specific assignment of resources (to be done). Note that it is suggested to remove WBS 2.2.2.8 and absorb the corresponding activity under WBS 2.2.2.1. New categories are "Performance Optimization & Monitoring" (suggested WBS 2.2.2.9), "MulitCore Support" (suggested WBS 2.2.2.10), and the already discussed "M&O, Testing" (suggested WBS 2.2.2.11). Assignment of resources to the new categories has not been finished yet and is expected to be rather need driven, i.e. dynamic. The initial activity for WBS 2.2.2.10 is funded by additional one-time support for FY10 and FY11 (0.5 FTE/year).

WBS	**Category**		**Comment**
	Old	**New**	
2.2.2.1	Framework	Framework	
2.2.2.2	EDM Infrastructure	EDM Infrastructure	
2.2.2.3	Detector Description	Detector Description	
2.2.2.4	Graphics	Graphics	
2.2.2.5	Analysis Tools	Analysis Tools	
2.2.2.6	Grid Integration	Grid Integration	
2.2.2.7	Core Service Usability	Core Service Usability	
~~2.2.2.8~~	~~Framework Update~~		now under 2.2.2.1
2.2.2.9		Performance Optimization & Monitoring	new category
2.2.2.10		MultiCore Support	new category
2.2.2.11		M&O, Testing	new category

4.1.2. WBS 2.2.3 – Data Management

At the time of the meeting we expected that "Distributed Data Management" (WBS 2.2.3.6) will be moved to a prospective new "Distributed Computing" WBS level 2 category. This may not be the case anymore. Then the only change to this particular WBS category is a change to a more descriptive title for 2.2.3.2, and the addition of the "M&O, Testing: category (suggested new WBS 2.2.3.8) – see table below.

WBS	Category		Comment
	Old	New	
2.2.3.1	Database Services & Servers	Database Services & Servers	
2.2.3.2	Common Data Management Software	Generic I/O Support & Performance	more specific title
2.2.3.3	Event Store	Event Store	
2.2.3.4	Non-Event Data Management	Non-Event Data Management	
2.2.3.5	Collections, Catalogues, Metadata	Collections, Catalogues, Metadata	
2.2.3.6	Distributed Data Management		moved?
2.2.3.7	Data Access Support	Data Access Support	
2.2.3.8		M&O, Testing	New

4.1.3. WBS 2.2.4 – Distributed Software

We expected this WBS category to be completely removed from the Software WBS and did not discuss this further.

4.1.4. WBS 2.2.5 – Application Software

So far WBS 2.2.5.6 (Application Support) has not been utilized. We suggest replacing it with "M&O, Testing", which collects all FTE efforts along these lines for WBS 2.2.5.1 through 2.2.5.5. This should give a better representation of the actual activity in the application software category. We also suggest to add a new category for "High Performance Detector Simulation", i.e. detector simulation faster than the standard ATLAS full simulation but with least significant loss of precision. Several US institutions are already involved here (SLAC, Pittsburgh,…), and we expect that this effort will ramp up in ATLAS to become more comparable to the detector simulation for physics in CMS, which is a factor of 10 faster than full simulation in ATLAS[11].

[11] Even when keeping in mind that the CMS detector is less complex than ATLAS, an improvement of that order for the massive ATLAS physics simulation, especially for backgrounds, is desirable.

WBS	Category		Comment
	Old	New	
2.2.5.1	Generator Support	Generator Support	
2.2.5.2	Tracking Infrastructure	Tracking Infrastructure	
2.2.5.3	Calorimeter Infrastructure	Calorimeter Infrastructure	
2.2.5.4	Muon Infrastructure	Muon Infrastructure	
2.2.5.5	Monitoring Infrastructure	Monitoring Infrastructure	
2.2.5.6	Application Support	M&O, Testing	more descriptive title
2.2.5.7		High Performance Detector Simulation	New

4.1.5. WBS 2.2.6 – Infrastructure

The infrastructure WBS does not need a "M&O" category, as we basically do not support any code development here. We suggest separating validation (now clearly indicated as "Software Validation") from software quality assurance, which now is part of the new suggested WBS category 2.2.6.3. This also reflects better the distribution of FTE across the activities, i.e. Nevski, Adams for validation, Ye and Undrus for librarian, and Undrus for software quality.

WBS	Category		Comment
	Old	New	
2.2.6.1	Quality Assurance & Validation	Software Validation	better description
2.2.6.2	Librarian	Librarian	
2.2.6.3		Software Quality Assurance & Infrastructure	New

4.1.6. WBS 2.2.7 – Analysis Software Support

This effort was moved from the software effort (WBS 2.2.) to its own level 2 category.

4.2. Planning of software activities

We like to improve the planning and documentation of software activities within US ATLAS. The concerns addressed here, besides the charge from the review committee, are continuity in transitions of personnel at all WBS levels, and establishing a reference to interesting software projects, thus attracting well motivated and competent people in case of replacements within, or extension of, the present workforce. In addition, the documentation should describe the actual tasks behind the WBS categories at two levels of detail: brief but meaningful and detailed. The more detailed level should include the description of actual tasks worked on, and future projects of interest, for each of the WBS level 4 categories. The brief descriptions should summarize the overall scope for each WBS level 3.

These descriptions are going to be provided on US ATLAS wiki pages, with the brief summaries collected on an introductory (top level) page linked from an appropriate referring page. The summaries should (again briefly) indicate actual priorities and expected future activities and kept at a level which allows non-software experts to understand the projects. This top level page then provides the links to one page for each of the WBS level 3 categories, which in turn contain the more detailed descriptions.

The detailed WBS level 3 pages should provide the actual activities and prioritized future plans at a level allowing software professionals even outside of ATLAS to understand them. Names, affiliations and very brief descriptions of tasks worked on should be provided for each individual.

Future plans include realistic expectations for activities for the next year (short-term, can be the next fiscal year) and the next 2-3 years (long-term, e.g. target up to FY12). Indications of how this planning fits into general ATLAS need to be provided. Priorities should be established considering continued, reduced, and enhanced future funding levels as guidelines. For this, no change in FTE and changes of ±1 FTE should be included in the planning.

The information content on these wiki pages needs to be reviewed and updated regularly. A review at the time of the quarterly reporting is appropriate. Updates should be considered each time a documented project is finished or significantly changed or extended in scope, and when a new project is added. The responsibility of maintaining these pages lies with the WBS level 2 manager for the top level page, and the WBS level 3 managers for the respective detailed documentation.

Action item: *PL to set up the top level page by the end of the ATLAS software week/mid September. L3 managers to set up the first individual pages by the end of September.*

4.3. Reporting

We feel that the reporting across the WBS level 2.2.X domains should be streamlined in format and content. The original idea developed at the retreat was to provide a list of clearly but briefly described deliverables for a given milestone date. The relevance of these deliverables with respect to the US ATLAS Software WBS levels 4 was to be clearly indicated. This scheme supports milestones for tasks across WBS categories. The suggestion for this format was dropped after Jim and Srini expressed concerns about its readability, especially for funding agents not being used to it.

Even when maintaining the present milestones per WBS category, a somewhat more common format for reporting across the WBS 2.2.X domains is desired. Some suggestions for a light weight format and some guidance for the content are given below.

4.3.1. Milestones

A milestone is defined within the scope of ATLAS software activities. The timeline is defined by the delivery dates for

- an Athena release (e.g., 16.0.0);
- a release of software not in, or only loosely coupled to, the Athena release cycle (e.g., external software, database releases,…);
- a software project like (re-)reconstruction and simulation projects with timelines not reflected by an Athena release;
- other appropriate deadlines associated with the US ATLAS software activity.

Milestones can extend beyond the next (closest) delivery date when appropriate. Milestones should be realistic within their respective timeline, multiple extension should be avoided as much as possible. It should be noted that the "Maintenance & Operations, Testing" WBS categories may not have milestones except if a certain software commitment is terminated.

4.3.2. Activity descriptions

The content granularity of milestones is given is guided by the WBS level 4 categories. A finer tracking of software efforts may be appropriate for certain activities.

The free text activity description for a given milestone should be brief but meaningful. Estimates for partial completion at a granularity of 20-25% should be given if appropriate. Milestones for WBS level 4, at least with the same timeline, should be grouped into more general milestones for level 3. These level 3 milestones in turn are collected into WBS level 2 milestones. A simple tabular format should be used to describe the milestones, their status and deadlines.

4.3.3. Effort tracking

US ATLAS management requested resource loaded effort tracking for each individual employed by the project. This issue was discussed during the retreat and a simple format is suggested (basically a matrix of personnel versus WBS category, filled with percentage FTE efforts, see example in Figure 1 on the right). It is suggested to submit the effort tracking with the quarterly reports. The data for this matrix will be collected by the WBS level 3 managers. Matrices can be merged by the WBS level 2 manager, if requested.

WBS / Name	2.X.Y.1	2.X.Y.2	2.X.Y.3	2.X.Y.4
AAA	100%			
BBB	25%		75%	
CCC		100%		
DDD	25%		25%	50%
EEE		50%		

Figure 1: Example for the effort tracking matrix described in the text.

Action item: *distribute spreadsheet templates for each WBS level 3 category by early September (PL).*

4.3.4. Reporting system

Activities related to M&O and testing can occupy a significant amount of FTE in times. Some of these can be best described as "continuing activity" without specific end date. We suggest implementation of the "no end date" task concept in the reporting system, or at least allow unlimited extension of milestones related to these tasks. In the latter case the date of the milestone would be the end of the fiscal year.

4.4. Change control

Several remarks from the review committee concerned the implementation of effort change control, which we interpreted as the reaction to requests from the overall ATLAS software effort for short or long term refocusing due to imminent needs. In general we like to avoid management overhead and the consequential slow processing of such requests, but we like some guidelines on how to react. One scenario discussed at the retreat introduces several levels of severity of these requests. Requests not involving changing or moving funding are:

Low level and very short term effort change within a given WBS level 4 category can be accepted by the corresponding level 3 manager if it has no significant impact on the efficiency of the effort as agreed upon in the MOU/sub-contract with the US ATLAS project office. Nevertheless this change needs to be reflected in the reporting, and notifying the level 2 manager is appreciated.

Temporary redirection across WBS level 4 categories within the same level 3 needs to be discussed with the WBS level 2 manager and presented by the level 3 manager in the regular phone conference with the US ATLAS computing management, or, if urgency requires, in a dedicated phone conference. Formal approval is not needed, but objections from level 2 and higher management should be addressed.

Long term/permanent redirection across WBS level 4 within the same level 3 or **any kind of redirection across WBS level 3** should be approved by the US ATLAS computing management or, if delegated accordingly, by the level 2 manager. The reasons need to be presented at the regular phone conference, or, again if urgency requires, at a dedicated phone call.

Any change requiring additional funds is understood to need approval by US ATLAS management. This includes temporary or long term relocations not covered by the approved funding for transitional periods of overlapping employments in case of replacement hires, etc.

5. Software topics

The discussion of software priorities was concentrated on the core (WBS 2.2.2) and data management (WBS 2.2.3) activities. The FTE assignments are estimates which in some cases have been produced after the retreat.

5.1. WBS 2.2.2 – core software

The following list is ordered by priority, starting with the highest first:

(1) **Memory crisis (1-1.5 FTE funded within target):** grid production limits the amount of virtual memory to 2 Gbytes/job for full reconstruction with Athena. This is presently addressed by the core software effort in US ATLAS. Part of the task is the development of appropriate tools monitoring the memory performance and finding problems. Temporary refocusing and redirection have started to support this effort.
(2) **MultiCore support (1 FTE from temporary funds):** new multi- and many core architectures are emerging fast. Adaptation of the reconstruction program for effective use of this hardware (e.g., event process parallelization) allows an effective use of these technologies. Both ATLAS and CMS

have received approval for dedicated funding for this effort, with ATLAS supporting 0.5 FTE in each FY10 and FY11. The developer has been identified.

(3) **Virtual machine support (0.5 FTE, funded within target):** the address limitations on 32 bit systems can be extended by installing several 32-bit virtual machines on one multi-core 64-bit node. This is in particular interesting for effective use of the Tier3 computing resources emerging in the US this and next year.

(4) **Adaptation to 64 bit architectures (0.5 FTE, funded within target):** study the performance of 64-bit architectures with respect to memory and cpu performance and optimize Athena for this architecture.

All high priority efforts except for the MulitCore software support, are covered from existing FTE, thus delaying some other lower priority work already assigned to the same personnel resources. The list above reflects the needs for initial data processing and the first year of ATLAS running.

5.2. WBS 2.2.3 – data management

The following list is ordered by priorities, starting with the highest first:

(1) **Non-event data support (no dedicated FTE assigned yet, expected to need 0.5-1.0 FTE):** this has so far been done at a 10% FTE level but now is a very essential effort for the initial experimental data and needs resources. These have to be provided by redirection and change of focus from existing resources in the data management domain. This must be addressed urgently, including estimates of the consequential delays in other efforts. It is notable that some of this effort is strongly linked to the computing facilities domain in US ATLAS.

(2) **ByteStream and POOL I/O optimization (0.5-1.0 FTE, funded within target):** US ATLAS provided help with ByteStream and general (POOL) I/O performance improvements. Read speeds are a major issue when the number of clients analyzing data rises as expected.

(3) **TAG database (0.5 FTE, funded within target):** TAGs are presently under-utilized, mostly because they are not needed in case of single physics source simulation analysis. With real data they become important for fast event selection and quick plotting of principal kinematic properties of reconstruction signal objects. Major problems are missing links from TAG to performance DPDs (sequence problem) and adding trigger information to TAGs (to be confirmed).

(3) **Metadata support (1.0 FTE, funded within target):** metadata support is important for experimental data analysis (e.g., good run lists). Supporting this effort may require some redirection and refocusing within the data management area. Details are emerging from the ATLAS wide metadata task force and include software infrastructure support for anomalous event tagging, good/bad run lists, general data quality feedback, book keeping, publishing of metadata, and debug streams. Eric Torrence agreed to consult with the developer(s) on the metadata content and structural aspects

As no additional resources are expected here, all priorities have to be addressed with existing FTEs through redirection or refocusing with existing efforts. It should also be noted that TAG database and metadata support have likely the same priority for real data when considering event picking. In any case, it was not obvious from the discussions at the retreat that one of those is more dominant than the other.

6. AOB

Annual in-person follow-up meetings are planned, with the possible extension of including analysis support, distributed computing, and possibly even facilities. The informal character should be maintained in any case. We also like to arrange quarterly dedicated phone conferences synchronized with the reporting dates. These are in addition to the weekly meetings and limited to WBS 2.2, with the possibility of extensions if needed.