

US ATLAS Tier-3 Whitepaper

Version 8
Aug. 8, 2006

The US ATLAS project has been asked to define the scope and role of Tier-3 resources (facilities or “centers”) within the existing ATLAS computing model and US ATLAS computing activities and facilities. This document attempts to address these questions by describing Tier-3 resources generally, and their relationship to the US ATLAS Software and Computing Project.

Originally the tiered computing model came out of MONARC (see <http://monarc.web.cern.ch/MONARC/>) work and was predicated upon the network being a scarce resource. In this model the tiered hierarchy ranged from the Tier-0 (CERN) down to the desktop or workstation (Tier 3). The focus on defining the roles of each tiered component has evolved with the initial emphasis on the Tier-0 (CERN) and Tier-1 (National centers) definition and roles. The various LHC projects, including ATLAS, then evolved the tiered hierarchy to include Tier-2’s (Regional centers) as part of their projects (Hoffman committee final report, CERN/LHCC/2001-004).

Tier-3’s, on the other hand, have (implicitly and sometime explicitly) been defined as whatever an institution could construct to support their Physics goals using institutional and otherwise leveraged resources and therefore have not been considered to be part of the official U.S. ATLAS Research Program computing resources nor under their control. We believe that this continues to be the case for Tier-3s, namely that *Tier-3s are not officially part of the US ATLAS Research Program*, meaning there is no formal MOU process to designate sites as Tier-3s and no formal control of the program over the Tier-3 resources. Tier-3’s are the responsibility of individual institutions to define, fund, deploy and support.

However, having noted this, we must also recognize that Tier-3’s must exist and will have implications for how our computing model should support US ATLAS physicists. Tier-3 users will want to access data and simulations and will want to enable their Tier-3 resources to support their analysis and simulation work. Tier 3’s are an important resource for U.S. physicists to analyze LHC data.

One important question is to what extent the Research Program should support Tier-3’s? For example, would we require that Tier-2 centers provide wide-area file-systems that Tier-3’s can access? What level of software “install” support could Tier-3 expect (if any)?

This document will define how Tier-3’s should best interact with the US ATLAS (and ATLAS) computing model, detail the conditions under which Tier-3s can expect some level of support and set reasonable expectations for the scope and support of US ATLAS Tier-3 sites.

Tier 3's in the ATLAS/US ATLAS Computing Model

The ATLAS computing model describes a hierarchical distributed virtual computing facility within which are defined Tier-1 and Tier-2 computing centers having certain specific MOU agreed roles and capacities to be used for the benefit and at the direction of ATLAS as a whole. The U.S. ATLAS Research Program management, together with international ATLAS, decides how these MOU pledged resources are used. This is accomplished in the U.S. Resource Allocation Committee (RAC)¹. In this model the primary functions of the Tier-1 are to host and provide long term storage for, access to and re-reconstruction of a subset of the ATLAS RAW data (20% in the case of the US Tier-1), provide access to ESD, AOD and TAG data sets and support the analysis of these data sets. The primary functions of the Tier-2's are simulation (they provide the bulk of simulation for ATLAS), calibration, chaotic analysis for a subset of analysis groups and hosting of AOD, TAG and some physics group samples.

US ATLAS has acted to establish compute capacity beyond the capacity it has pledged to meet the obligation of international ATLAS to be used specifically for the benefit US ATLAS physicists. This US ATLAS specific computing is located at the Tier-1 and Tier-2's making use of the infrastructure and operational expertise required there anyway, at a scale of ~50% (for the Tier-1) of the level of the capacity being pledged to international ATLAS. US ATLAS decides how these resources are used by means of the Resource Allocation Committee, not the local Tier-1 or Tier -2's or international ATLAS.

Tier-3 sites are institution-level non-ATLAS or US ATLAS funded or controlled centers/clusters which wish to participate in ATLAS computing, presumably most frequently in support of the particular interests of local physicists (physicists at the local Tier-3 decide how these resources are used). These are clusters of computers which can vary widely in size. It should be noted that substantial institutional funding to originate such clusters is potentially available, and that they could make a real contribution to the impact of US ATLAS on the overall ATLAS physics output. As such, there is considerable value in providing some level of technical support to these sites.

Support issues (financial, technical expertise, services)

- Individual ATLAS institutions are expected, out of their local resources, to buy individual physicist's equipment, laptops, desktops, printers, etc.
- An individual physicist's share of the ATLAS and US ATLAS resources (at Tier-1 and Tier-2's) in combination with modest local computing resources (which could be just a modern desktop machine for each physicist) should be sufficient to accomplish required computing tasks for ATLAS and for effective participation in physics analysis.

1

- The Tier-1 and Tier-2's have as primary responsibilities to support such analysis by their users with capacity shares and priorities being established by the RAC for US ATLAS controlled resources together with international ATLAS management for the resources pledged to ATLAS as a whole.
- Sites having significant institutional or base grant-funded computing centers or clusters are encouraged to use them for analysis or other ATLAS computing activities.
- Support from the Tier-1 and Tier-2's to such Tier-3 centers in terms of expertise (install, configure, tune, troubleshooting of ATLAS releases and the OSG stack) and services (data storage, data serving, etc.) follows from responsibility to support the US ATLAS user community. This support would have to be limited to Tier 3 sites with standard ATLAS operating systems.
- Larger Tier-3 sites should be or should become participants in OSG and so get additional technical support via that path.

Part of our task is to set reasonable expectations for the size and scope of Tier-3 centers. We recognize that there will likely be extremely large variances in the amount of computing power and storage at US ATLAS Tier-3 sites. One could reasonably define a Tier-3 as anything a US ATLAS institution so designates, larger than a single machine. We fully expect that some Tier-3 sites may have resources to rival a Tier-2 (or perhaps even the Tier-1!). Our goal is not to constrain the definition of a Tier-3 but to determine a reasonable capability for a Tier-3.

The typical scaling from the MONARC model was to assume that the Tier-0 would provide about 1/3 of the total resources for an LHC project and the integrated Tier-1's would provide about 1/3 with the last 1/3 provided by the integrated power of the global Tier-2's. In the US ATLAS case this implied that the five Tier-2's would each contribute roughly 1/5 of the Tier-1. Although Tier-3's may be any size, we expect most of them to be smaller than a Tier-2.

Alternatively we could estimate a suggested Tier-3 capacity by determining the type of activities a Tier-3 would be expected to support and scale accordingly. This is perhaps the best means of determining what a "typical" Tier-3 requires in computing power, network connectivity and storage.

We envision the following to be typical examples of uses of a Tier 3:

- Interactive analysis of Ntuples. This requires no direct connection to the ESD or AOD, but it does require access to the data when these Ntuples are generated.
- Development of analysis code. This would motivate a local copy of a small number (perhaps a few thousand) of ESD, AOD, or RAW events. It would be desirable for at least some fraction of these events to be complete "vertical slices"

- having the RAW, ESD, AOD and TAG for the same events.
- Running small local test jobs before submitting larger jobs to the Tier-1 or Tier-2 via the grid. This would motivate similar sized copies of the data as above. It also motivates having access to at least the appropriate subset of the TAGs at the Tier-3, because this is the same selection mechanism that will be used when the full scale job is run,
- Running skimming jobs of the Tier-1 and Tier-2's via the grid, and copying the skimmed AOD (or rarely ESD) back to the Tier-3 for further analysis. The output of this skim must be a very small subset of the AOD – of order a few percent.
- Analyzing the above skimmed data via Athena.
- Production of MC samples of special interest to the local institution.
- For larger Tier-3 centers, opening those resources to ATLAS managed production as well as individual ATLAS physicists via OSG Grid interfaces and the ATLAS VO authentication, authorization and accounting infrastructure. Guidance for establishing policies for queue priorities and/or storage may be discussed in the RAC.

These use cases can be met by large or small clusters at Tier-3 centers with the standard OSG software suite installed as well as ATLAS releases, the ATLAS Distributed Data Management end user tools (DQ2), and potentially TAG databases or files. This is a well established process at the U.S. Tier-1 and Tier-2 sites (though some problems are still being worked out) and we expect that support for installing these software suites will be the extent of U.S. Research Program support at Tier-3 centers.

Summary

- Some local compute resources, beyond Tier-1 and Tier-2, are required to do physics analysis in ATLAS.
- These resources are termed Tier-3 and could be as small as a modern desktop computer on each physicist's desk, or as large as Linux farm, perhaps operated as part of a shared facility from an institution's own resources.
- Resources outside of the U.S. ATLAS Research Program are sometimes available for Tier-3 centers. A small amount of HEP Core Program money can sometimes leverage a large amount of other funding for Tier-3 centers. Decisions on when it is useful to spend Core money in this way will have to be considered on a case by case basis.
- Support for Tier-3 centers can be accommodated in the U.S. Research Program provided the Tier-3 centers are part of the Open Science Grid and that they provide access those resources with appropriate priority settings to US ATLAS via the VO authentication, authorization and accounting infrastructure.