

ATLAS GREAT LAKES TIER-2 RENEWAL PROPOSAL

AGLT2 History and Status

The ATLAS Great Lakes Tier-2 started operations in September 2006 after successfully competing to become one of five US ATLAS Tier-2 centers. We were successful in large part due to the strong support we received from the University of Michigan and Michigan State University (our Tier-2 center is composed of sites at both Universities). This support has helped enable our site to become one of the top LHC (Large Hadron Collider) Tier-2s in the world.

Shown in Table 1 are the resource data from the last 5-year period for AGLT2. We have managed to deliver more than we both proposed and agreed to in our MOU by carefully allocating our resources where they would bring the most impact.

	2006-07	2007-08	2008-09	2009-10	2010-11
AGLT2 Original Proposal					
CPU(HS06)	3348	5000	7792	10972	
Storage(TB)	304	471	770	1055	
ATLAS MOU Values for AGLT2					
CPU (HS06)	2324	3860	4960	11040	
Storage(TB)	155	322	520	1060	
Delivered Resources at AGLT2					
Job-slots	650	1140	2300	4244	5477
CPU(HS06)	4840	9050	19800	35648	46557
Useable Storage (TB)	170	662	1139	1927	2839

Table 1 Prior 5-year period resources from the last proposal, the MOU amount and the delivered amount for AGLT2.

We should note that our Tier-2 has been one of the top Tier-2 WLCG wide since accurate accounting has been kept starting in January 2009. Figure 1 shows the normalized delivered CPU-hours by “Tier-2” for all of WLCG from January 2009-September 2010. AGLT2 was second overall; the first Tier-2 represents a set of 11 Italian sites that jointly support all four of the main LHC VOs: ALICE, ATLAS, CMS and LHCb. During the almost 2 year period of accounting results AGLT2 was frequently the top site worldwide during many months.

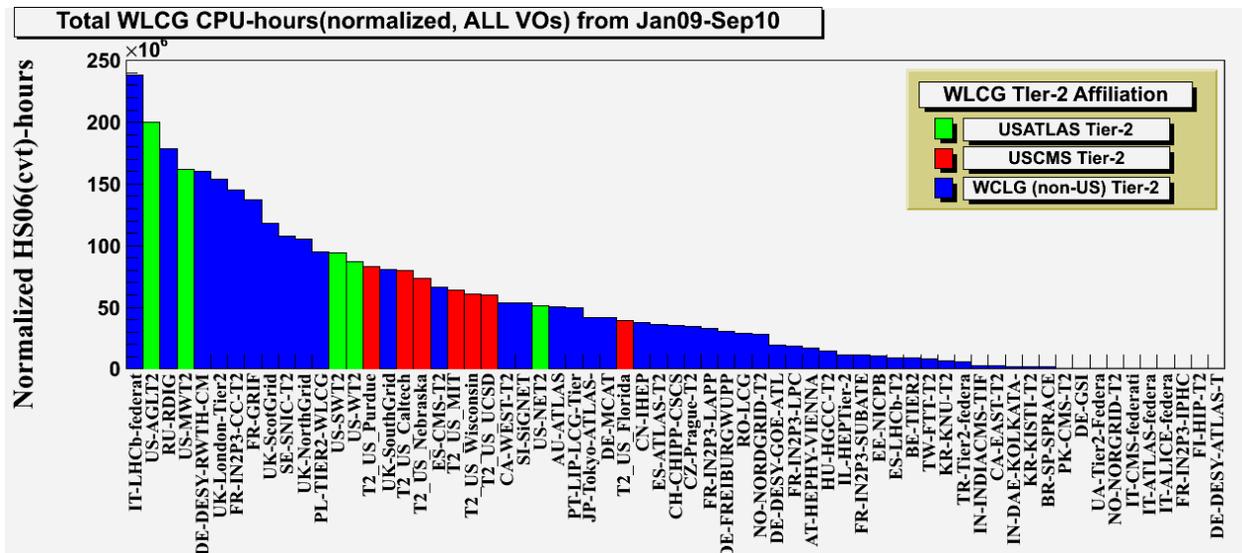


Figure 1 WLCG accounting for delivered CPU-hours for all WLCG Tier-2s from January 2009 through September 2010

Infrastructure Details for AGLT2

The AGLT2 is currently hosted at two sites: the University of Michigan and Michigan State University. Both sites have provided power, space and cooling for the Tier-2 facilities during the last 5 years and will continue to provide such support for the next 5 years. This is a significant contribution to the project and is estimated at approximately \$180K/year over the two sites.

The University of Michigan site hosts 8 racks for equipment, 3 racks hosting an APC Symmetra 80kW UPS system and 1 network rack. The UPS system is backed by an external generator which can switch on in 30 seconds in the event of a power failure. Current UPS battery capacity is 2 minutes at full load. The machine room has two Liebert air-conditioning systems, each capable of providing all required cooling. In addition the Tier-2 has access to the Tier-3 space in the physics department which uses water cooled racks as well as building air-conditioning to support up to 6 racks of space. Available power in excess of 60kW is wired into overhead power trays. A dedicated fiber (currently running at 10GE) connects this room with the Tier-2. Currently the Tier-2 has migrated some out-of-warranty systems to this space so that we can continue to operate them for supporting Tier-2 work.

The Michigan State University location was originally renovated to host AGLT2 at the beginning of the past 5 year period. It can host up to 15 racks of Tier-2 equipment. The room has 225KW central PDU, 1 30 ton Liebert CRAC, 1 22 ton Liebert CRAC. Total cooling is 182kW with 120kW power/cooling for Tier2 use. Also, the power is generator backed. Worker nodes installed at MSU don't have UPS power protection but all storage and server elements are installed with rack-level UPS systems

Networking Infrastructure for AGLT2

Both sites benefit from the excellent network infrastructure our Universities have put into place. In Figure 1 below is a diagram of the fiber path comprising MiLR (Michigan Lambda-Rail) which serves AGLT2. Each path has 3 fiber pairs and each pair can support multiple 10GE wavelengths. Current hardware supports up to 8 wavelengths and up to 40 total wavelengths can be supported by adding additional cards (8 channels/card; 5 cards maximum per chassis). We

also anticipate a future upgrade to MiLR (going to 40 Gbps/wave or 100Gbps/wave) sometime during the next 5 years.

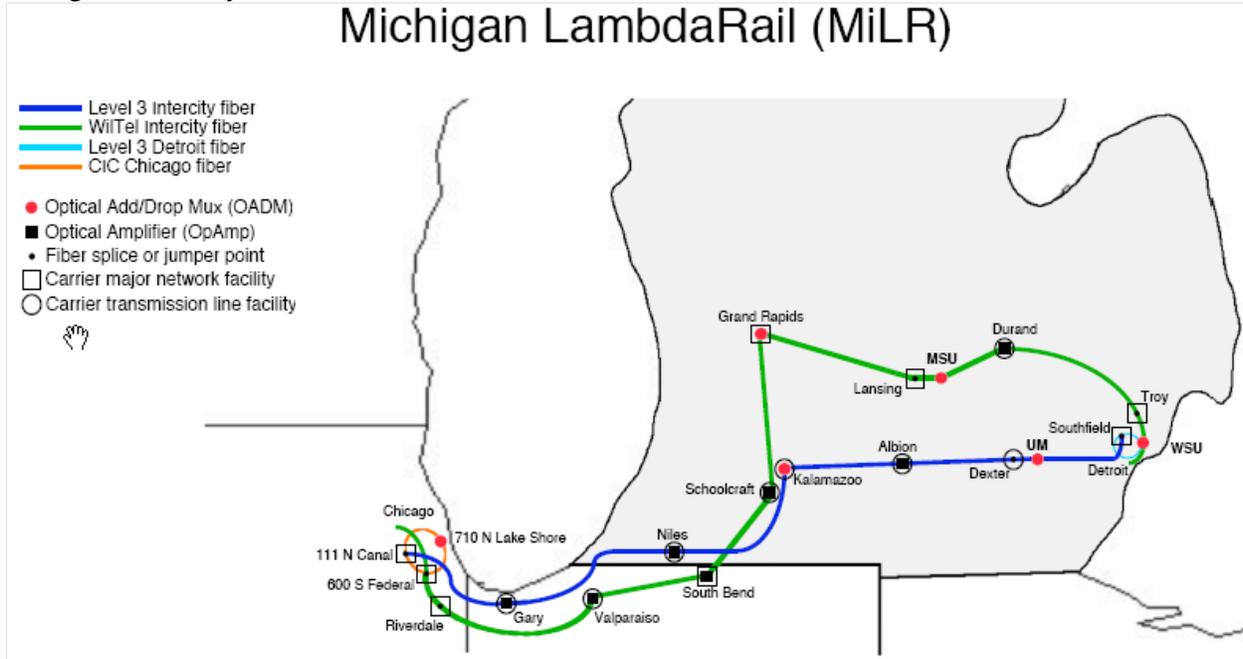


Figure 2 MiLR fiber paths used by AGLT2

Currently the Tier-2 has dedicated wavelengths from each site to the Chicago peering point and a shared wavelength between the sites via the Detroit direction. The MSU wavelength costs will be covered by MSU for the next 5 year period. At UM we have agreed in principle that the project should not have to bear the costs for the external 10GE network connectivity to Chicago and to MSU, though the details are being worked out still. At UM we also have dedicated use of a second wavelength to Chicago for research purposes and Tier-2 backup. This extra 10GE wavelength may not be available for the full duration of the next 5 years. With the MiLR infrastructure we have the ability to quickly and easily expand our network capacity if demand calls for it.

In addition to the WAN charges we also need to provide \$9K/year of peering charges for our connections to the rest of the world. This will need to be supported by the project.

Resource Division by Site

We currently divide the \$600K/year into \$375K for UM and \$225K for MSU and this fractional split is planned to continue for the next 5 years. It is possible this proportion could change if circumstances dictated.

Resource Projections for AGLT2

In providing the resource projections we use the following table of costs to determine the impact of future equipment funds on computing and storage resource capacities. Both HS06 and TB assume a 3-year doubling time (factor of 1.26 per year). That is for a fixed cost the specified quantity increase by a factor of 1.26 every year. The numbers used are based upon “Tier-2”

suitable equipment as available from Dell using their LHC pricing as of October 2010. These numbers are the agreed upon ones with USATLAS.

Cost Projections	FY12	FY13	FY14	FY15	FY16
\$/HS06 (CPU power)	\$14.29	\$11.34	\$9.00	\$7.14	\$5.67
\$/TB (useable)	\$216.68	\$171.98	\$136.50	\$108.34	\$85.99

Table 2 Costs by year based upon \$273/TB (useable; powerful server) and \$18/HS06 in fall 2010

We present the information for the next 5 year period via 3 tables showing the funds available to purchase equipment (further broken down by storage and compute), a table of the corresponding delivered resources (including the amount added and removed each year), and finally a table of the costs in each category as a percentage of the total available resources and separately for project salary and equipment as a fraction of the projects funds (\$600K/year; escalated by 1.5% each year starting in FY13).

AGLT2 Resources Delivered

The following tables provide the details about AGLT2 resources for the next 5 year period.

AGLT2 by Year	FY12	FY13	FY14	FY15	FY16
Equipment Budget	\$211,984	\$225,640	\$225,407	\$225,031	\$224,504
Compute Budget	\$84,794	\$90,256	\$90,163	\$90,012	\$89,802
Storage Budget	\$127,191	\$135,384	\$135,244	\$135,018	\$134,702

Table 3 Available equipment budget by year for AGLT2

Table 3 summarizes the available equipment budget (after paying for personnel, travel, networking and parts). We assume a 40 to 60% split between compute and storage which gives us the best match to the ATLAS requirements.

Year	FY12	FY13	FY14	FY15	FY16
Job Slots	5177	4867	4219	3870	4970
HS06	46596	43805	37975	34827	44731
TB(useable)	2711	3021	3224	3388	4368
HS06 added	5935	7960	10018	12601	15839
HS06 removed	4210	10750	15848	15749	5935
TB added	587	787	991	1246	1566
TB removed	492	477	788	1082	587

Table 4 Delivered resources for AGLT2 with detail on added and removed resources

Note that for the delivered resources show in Table 4, we are assuming we always are removing resources from the 5th year prior, e.g., in FY16 we are removing resources purchased in FY12. The amount of TB and HS06 added and removed are shown in the bottom of the table.

The following table breaks down AGLT2 expenses estimated for the next 5 year funding period using the total available resources (local and project) for AGLT2 in the top part of the table. The corresponding FTEs associated with AGLT2 (both local and project) are shown highlighted in yellow.

AGLT2 Cost Percentages	<u>FY12</u>	<u>FY13</u>	<u>FY14</u>	<u>FY15</u>	<u>FY16</u>
Equipment	20.4%	21.4%	21.1%	20.8%	20.4%
Personnel (project)	18.9%	19.2%	19.5%	19.8%	20.1%
Personnel (non-project)	21.5%	21.8%	22.1%	22.5%	22.8%
Power/space/cooling	17.3%	17.1%	16.8%	16.6%	16.4%
Overhead	12.7%	12.8%	12.9%	13.1%	13.2%
Networking	9.4%	9.3%	9.1%	9.0%	8.9%
Parts/supplies	1.0%	0.9%	0.9%	0.9%	0.9%
Travel	1.4%	1.4%	1.4%	1.4%	1.4%
FTEs AGLT2 total	5.0	5.0	5.0	5.0	5.0
Proj FTE	1.9	1.9	1.9	1.9	1.9
Other FTE	3.1	3.1	3.1	3.1	3.1
%Project Equipment	35.3%	37.1%	36.5%	35.9%	35.3%
%Project Salary	32.8%	33.2%	33.7%	34.2%	34.7%

Table 5 AGLT2 costs as a fraction of total resources. The bottom two rows show equipment and project personnel costs as a fraction of only project funds.

As noted, the fractional breakdown information shown in Table 5 (above the “FTEs AGLT2 Total” row) are in terms of total available resources, both project and non-project. The bottom two lines of the table show the fraction of project funds that are spent on equipment and personnel.

Intangibles and Further Comments for AGLT2

During the first 5 years of USATLAS Tier-2s, the AGLT2 has contributed to ATLAS in numerous ways. Specifically, the following list highlights some of the infrastructure areas where we have made significant contributions:

- Reliable system configuration and system optimization
- Network research supporting data-intensive WAN flows
- Throughput monitoring and testing
- Storage system research and benchmarking
- Grid, cluster and network monitoring
- Configuration management and node provisioning

We view our role as a Tier-2 to be much more than a simple service provider to (US)ATLAS but rather an active participant in researching, testing and enabling cost-effective infrastructures capable of effectively supporting ATLAS physicists and, more generally, all those involved in data-intensive, distributed science. We believe we have demonstrated that we are a very reliable and effective resource for (US)ATLAS and will be able to continue to be so for the foreseeable future.