1 Tier 2 and Tier 3 Facilities

Over the past several years our group has played a leadership role in the creation of university-based computing infrastructure for the US ATLAS physics and computing program and for data-intensive, distributed science generally through Open Science Grid. We have deployed a high performance computing facility (the ATLAS Midwest Tier 2 Center) and we manage the overall facilities integration program for US ATLAS as part of the Physics, Software and Computing project. We leverage these activities and resulting knowledge base to support a Tier 3 computing resource for UC physicists.

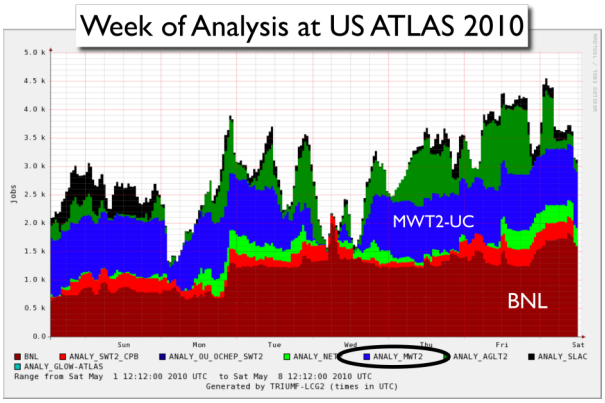

The Tier 2 facility is a joint operation with Indiana University with overall leadership and the majority of resources at UC. The facility provides 2352 CPU-cores and 1.41 Petabytes of disk space for the ATLAS collaboration. 519 ATLAS users have submitted analysis jobs via the Grid to the center (or roughly half of all active ATLAS users). Currently the maximum number of analysis jobs is 1200 with the remaining job slots being backfilled by Monte Carlo production tasks using priorities established by the Collaboration and the US ATLAS Resource Allocation Committee. The aggregate computing power is rated at 17,500 "HEP Spec 2006" units, which is the current CPU benchmark adopted by CERN as a way of comparing resources contributed to the Worldwide LHC Computing Grid (WLCG). This is roughly 30% of that provided by Brookhaven National Lab (the US ATLAS Tier 1 center). US ATLAS provides roughly 30% of the overall computing capacity for international ATLAS Collaboration.

At UChicago the Tier 2 center physically consists of 257 servers, 1528 CPUs, 1,394 disk drives, two managed Ethernet switches, and an uninterruptable power supply system all housed in 17 racks in the basement of the Research Institutes building on campus. 225 of the servers are for compute nodes, the remainder are for storage and Grid service management. Networked power distribution units and serial console servers are used to manage the systems. The storage servers have 10 Gigabit/s network cards for high analysis bandwidth. We've optimized data access to these servers so that ATLAS analysis jobs are more than 70% efficient in terms of CPU versus wall-time (a measure of IO efficiency), making it one of the most efficient sites in ATLAS. Wide area transfers are achieved through a fiber optic uplink from our center to a major research network peering hub in Chicago. With ESnet we established a direct-path 10 Gigabit/s circuit to Brookhaven National Lab. We have studied and tuned this network extensively and can achieve sustained 1000 MByte/sec transfers as needed for fast downloads ATLAS datasets into the center. The figure above illustrates high bandwidth transfers over a 24-hour period in April 2010. These data were downloaded to Chicago while simultaneously running analysis jobs, indicating good stability of the required services.

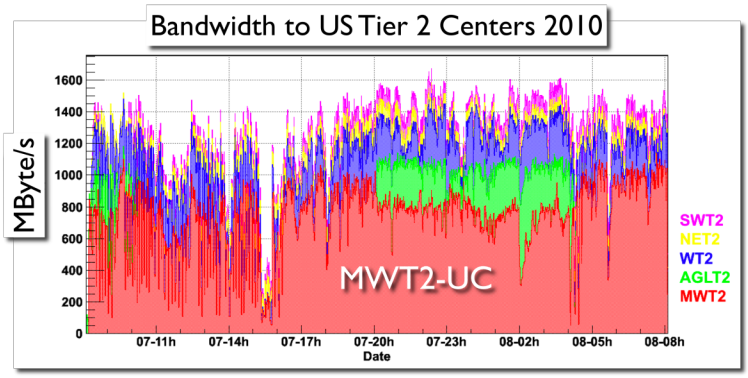

The operational reliability of the center (job efficiency, data transfer efficiency and other functional tests) as measured by external ATLAS probes has consistently exceeded 95%, with service provided 24x7x365

except for maintenance downtimes. The computing technical staff is funded at the level of 1.33 FTE from the NSF Tier 2 grant (PHY 06-12811), with the University providing the remaining 0.67 FTE making a team of two Unix systems administrators. These administrators work closely with the ATLAS physics and computing community as necessary to troubleshoot job and data management issues. The University also provides machine room space, power and cooling (180 KW, 55 Tons).

The Tier 2 center hosts datasets of various formats as needed for physics analysis, including Monte Carlo, reconstructed data and user analysis datasets. The figure below shows the dramatic surge in hosted data (in Terabytes) at UC from October 2008 to July 2010 covering the period from cosmic running to colliding beams at the LHC. Not shown is an equal amount Monte Carlo data. User data accounts for about 200 Terabytes which users download to their home institutions when their jobs finish. Our data catalog includes some 5.7M files in over 61K datasets. The storage capacity of the Tier 2 will increase to 2.5 Petabytes during the next 12 months to accommodate the increasing ATLAS storage requirements.

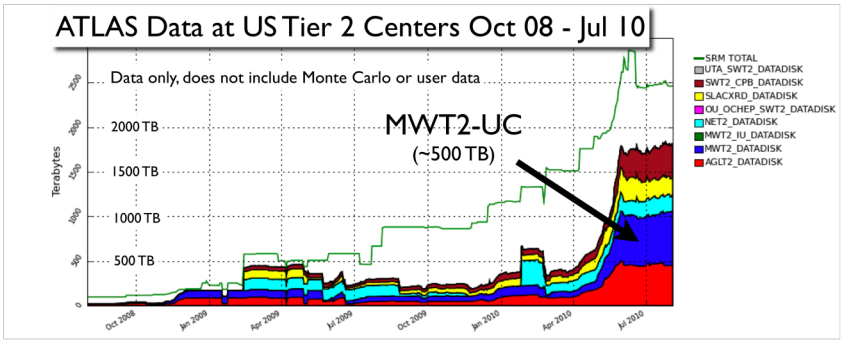

Our Tier 3 facility consists of 28 commodity servers procured four years ago with University funds. It shares machine room space with the Tier 2 and provides 112 CPU-cores and 40 Terabytes of storage for local batch and interactive analysis. It is integrated with the ATLAS data management system to enable managed transfers and cataloging of datasets. Where possible we re-use cluster and data management tools developed for the Tier 2. We plan to upgrade the cluster this year using $63K from a multi-institution NSF-MRI grant *MRI-R2 Consortium: Development of the U.S. ATLAS Physics Analysis Instrument (APAI) for the Analysis of Data from the ATLAS Experiment at the Large Hadron Collider*.