

# The Neutron Science TeraGrid Gateway, a TeraGrid Science Gateway to Support the Spallation Neutron Source

John W. Cobb<sup>\*</sup>, Al Geist<sup>\*</sup>, James A. Kohl<sup>\*</sup>, Stephen D. Miller<sup>‡</sup>, Peter F. Peterson<sup>‡</sup>, Gregory G. Pike<sup>\*</sup>, Michael A. Reuter<sup>‡</sup>, Tom Swain<sup>†</sup>, Sudharshan S. Vazhkudai<sup>\*</sup>, Nithya N. Vijayakumar<sup>§</sup>

## SUMMARY

The National Science Foundation's (NSF's) Extensible Terascale Facility (ETF), or TeraGrid [1] is entering its operational phase. An ETF science gateway effort is the Neutron Science TeraGrid Gateway (NSTG.) The Oak Ridge National Laboratory (ORNL) resource provider effort (ORNL-RP) during construction and now in operations is bridging a large scale experimental community and the TeraGrid as a large-scale national cyberinfrastructure. Of particular emphasis is collaboration with the Spallation Neutron Source (SNS) at ORNL. The U.S. Department of Energy's (USDOE's) SNS [2] at ORNL will be commissioned in spring of 2006 as the world's brightest source of neutrons. Neutron science users can run experiments; generate datasets; perform data reduction, analysis, visualize results; collaborate with remotes users; and archive long term data in repositories with curation services. The ORNL-RP and the SNS data analysis group have spent 18 months developing and exploring user requirements, including the creation of prototypical services such as facility portal, data, and application execution services. We describe results from these efforts and discuss implications for science gateway creation. Finally, we show incorporation into implementation planning for the NSTG and SNS architectures. The plan is for a primarily portal-based user interaction supported by a service oriented architecture for functional implementation.

KEY WORDS: Portal, Neutron Scattering, TeraGrid, Science Gateway, Service Architecture, Grid

## 1. INTRODUCTION

**Neutron Science:** Neutron scattering is used to determine the structure, dynamics, order, and magnetic response of matter in diverse disciplines such as materials science, biological science, nanotechnology, chemistry, engineering mechanics, polymer physics, earth science, and fundamental physics [3]. As a diagnostic tool, neutron scattering provides unique information about material properties including positions of atomic nuclei and characterization of magnetic properties. Neutron scattering is also useful when used in association with other micro- and nano-diagnostics such as synchrotron light sources, atomic force microscopy and electron microscopy. The neutron science community consists of a few thousand users who, collectively, conduct on the order of 10,000 experiments annually at the five facilities in the U.S. and a few dozen worldwide. The users' current main focus is on experimental conduct during beam time and afterward on data collection, reduction,

---

<sup>‡</sup> Spallation Neutron Source, Oak Ridge National Laboratory, {millersd, petersonpf, reuterma}@ornl.gov

<sup>†</sup> Department of Computer Science, The University of Tennessee, swain@cs.utk.edu

<sup>\*</sup> Computer Science and Mathematics Division, Oak Ridge National Laboratory, {cobbjw, gst, kohlja, pikeg, vazhkudaiss}@ornl.gov

<sup>§</sup> Department of Computer Science, Indiana University, nvijayak@cs.indiana.edu

and analysis. This sequence currently occurs in a fairly serial fashion. Most of the results are not known until well after the experimental access window has closed.

The community is currently at a tipping point because of two major events. First is the impending commissioning of next generation pulsed sources (SNS at ORNL in 2006 [2], J-PARC in Japan around 2008 [4], and the ISIS second target station project (TS2) in 2007 [5]) which are each at least an order of magnitude brighter and will create several orders of magnitude larger data sets. Second, the community is beginning to take advantage of opportunities for the use of advanced software and cyberinfrastructure for experimentation, data acquisition, data reduction, data analysis, and also some forms of remote presence and even remote control. The SNS is of particular interest since it is a close partner with the TeraGrid's ORNL-RP effort, the NSTG, and because it is the first next generation facility to commence operations. A deeper exploration of the requirements of SNS' user base for data reduction and analysis has been conducted. Those results have been used to shape the NSTG's requirements and plans.

The data production rates, computational, short-term storage and archival storage requirements are all exploding by factors of 10 to 1000. New, more powerful sources will, necessarily, cause changes in how to use data to drive scientific discovery. Future scientific breakthroughs may hinge on efficient and rapid data analysis and management and sophisticated software. If effective capability expands by a factor of a 1000 then users can perform 100 times more experiments each with 10 times more data. Alternately, the factor of 1000 could be spent to create experiments with 1000 times more data for the same number of experiments. These prospects will create stresses on responsiveness of data reduction and analysis infrastructures. It will test the scalability of software and cyber infrastructures. Today, we can permanently affect the development course of facilities costing billions of U.S. dollars with 40 year lifetimes! However, two challenges must be faced. First neutron sources of the future must handle data differently. Second facility construction budgets often do not include resources to completely capitalize on these opportunities. In response the international community initiated several collaboration opportunities, especially the Neutron Science Software Initiatives (NeSSI) workshops series, the results of which are described in more details in section 4.

**Extensible Terascale Facility, TeraGrid:** The NSF has sponsored the construction and now the operation of the TeraGrid [6]. It is the largest national cyberinfrastructure project. The TeraGrid now consists of eight different sites offering many large scale and varied computing, data, networking, and other services including more than 60 TeraFlops of computing, and over 1 PetaByte of storage all connected by one or more 10 Gigabit network connections. Additional special resources include visualization hardware, instrument services, data collections and database services. Currently the ETF has dual thrust areas of "Deep" and "Wide". The "Deep" efforts provide the traditional supercomputer center services of highest end computational resources for cutting edge computational science. The "Wide" focus aims to enlarge the community of cyberinfrastructure users by engaging non-traditional computational scientists through new and innovative means. The TeraGrid science gateways effort and the ORNL-RP are both directed toward this "wide" focus.

**Neutron Science TeraGrid Gateway (NSTG):** The ORNL-RP effort has, from the start of construction in 2003, been narrowly focused on creating a bridge between the ETF cyberinfrastructure and the neutron science community in general and the SNS in particular. This effort has been dubbed the Neutron Science TeraGrid Gateway or NSTG. While the TeraGrid Science Gateway effort has

become a focus during the transition to operations, the NSTG has been central to the ORNL-RP effort as far back as 2003.

The NSTG is assisting the neutron science community to adapt to the new data and computational challenges at next generation sources. Traditionally at the end of experimental runs, users depart with their raw datasets. Afterwards, at home institutions, datasets are analyzed using fitting and visualization tools put together by themselves or rely on existing, third-party applications such as MATLAB [7], DAVE [8], or ISAW [9]. Although those software tools are a significant advance over previous tools, there are disadvantages. First, each user is alone and isolated in terms of the reduction/analysis and visualization tools. Second, users will not have the necessary computing and storage infrastructure at their home institutions for dataset sizes from new sources. Scientific insights, today, are almost always driven by computationally intensive analyses of massive datasets. Third, since, analysis results are derived from custom solutions by users, working independently of each other; there is little validation of results and limited ability to share results. The ultimate impact is that opportunities to accelerate the publication cycle are lost. Users currently take almost a year-and-half from when they conducted their experiments to generate quality publications.

With the construction of the SNS the desire has emerged for the facility to act as a portal for users to look to for all of their needs. Recent technological trends make this possible. First, the notion of science gateways and portals is emerging as an elegant way to provide easy Web-based as well as programmatic access to a host of complex, backend services including analysis workflow executions, high-end computing, data services, visualization capabilities, metadata searches, etc. Gateways can hide the complexities involved in seamlessly accessing remote services. Second, the advent of the TeraGrid and ORNL's involvement as a resource provider implies that the SNS user base can potentially gain transparent access to several teraflops of computing made available through a grid infrastructure.

In this paper, we present the in process design and implementation of the NSTG and discuss our experiences and user feedback based on a prototype implementation. Section 2 discusses the SNS facility and some of its requirements in modest detail. Section 3 addresses related work in other TeraGrid gateway efforts, portal technologies, and frameworks. In Section 4, we present a summary of the NSTG and SNS SW architecture requirements gathered from extensive surveys of user groups. In Section 5, we present the design and prototype implementations of the various pieces of our services architecture. Section 6 presents a discussion based on our prototype and future plans.

## **2. SPALLATION NEUTRON SCIENCE (SNS) BACKGROUND**

The U.S. Department of Energy's (USDOE's) SNS, is scheduled for completion in June of 2006. As of April 2005, the 1.4 billion \$US construction project was 92% complete. The current state of the site in April of 2005 is shown in Figure 1. The SNS is an accelerator-based pulsed neutron source. The accelerator complex will deliver 1.4 megawatts average power (23.3 KJoule pulses at 60 Hz). The SNS target is composed of a closed loop of liquid mercury. Through spallation reactions, the incoming protons react with the mercury, to create copious neutrons. 18 beam ports (some of which can be used for multiple instruments) emanate from the 3 target moderators. Out of 24 possible beam lines 17 have been allocated and 16 have full or partial construction funding. The current instrument layout is shown in Figure 2. Figure 3 shows the instrument commissioning profile over time. The red triangles show

the number of commissioned instruments over time. The blue diamonds show the planned integrated intensity (power X hours of user time). The blue square in the lower left corner denotes the current integrated intensity of the ISIS source in the UK, currently the world's most powerful pulsed neutron source. In addition to this current configuration, the USDOE is already considering SNS upgrade proposals for 2 MWatts of accelerator beam power and a second target station, potentially with another 24 beam lines.

The SNS is a pulsed source. The instruments are time of flight based. Neutrons of many different energies and wavelengths are created almost simultaneously in the target monolith. As the neutrons propagate outward toward instrument sample areas, the neutrons separate according to energy. The energy of a neutron is discerned by the time taken to travel to the sample. Each SNS experimental instrument has several detectors that can detect and locate scattered neutrons both as to position (scattered angles) and timing. Raw data can be given in two forms: event mode where each neutron detected is recorded based on its time and detector position; and histogram where the data is binned into an abbreviated list of position and time channels and the sum of counts falling within the bin boundaries are recorded. The raw data rates are then expressed as proportional to neutron counts or to histogram size.

There are 60 pulses per second. The maximum data rate is given as a product of data size per pulse times the repetition rate. However, for many experiments, data will be summed over many pulses from minutes to hours. The final result will be either an aggregate list of neutron events (event mode) or a single integrated histogram.

The SNS will deliver neutron intensities eight (8) to one hundred (100) times higher than current sources.

Further, neutron beam line and instrument technology improvements such as super-mirror guides, non-focusing optics, and increased solid angle detector coverage imply another increase of ten (10) to one hundred (100). There will be an equally dramatic increase in the size of the resulting data sets. The challenge is to translate this capability and data growth into significant scientific results. In addition to



Figure 1: View of the SNS, April 2005

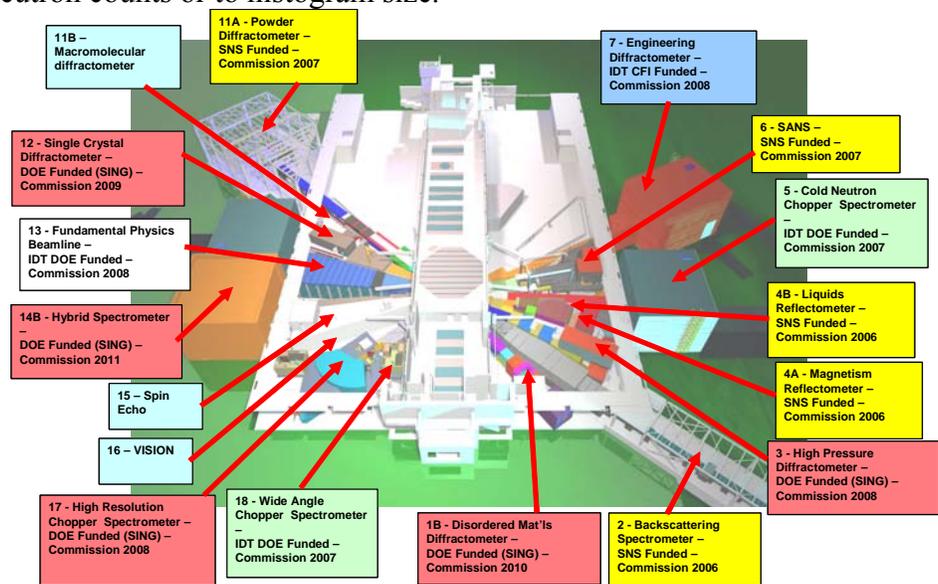


Figure 2: SNS Instrument Layout

raw data, intermediate data may be captured and stored for some length of time. Intermediate data can be used for real-time interaction to provide ongoing diagnosis of experimental progress. Finally, raw data may be captured in very rapid succession, perhaps even close to on a per pulse basis when some physical parameter of the sample is changing over time such as sample temperature, magnetic field, or even sample position.

The first step in data management is capture of needed raw (and perhaps intermediate) data. Next is reduction. Typically reduction involves the routine workflows to convert detector counts and timing into physical meaningful quantities such as a scattering function as a function of momentum transfer,  $S(q)$ , or momentum and energy transfer  $S(q, \omega)$ . Reduced data sets are much smaller than raw data. Therefore, the primary data challenge is raw data. Annotation and provenance are the challenge for reduced data and further downstream analyses.

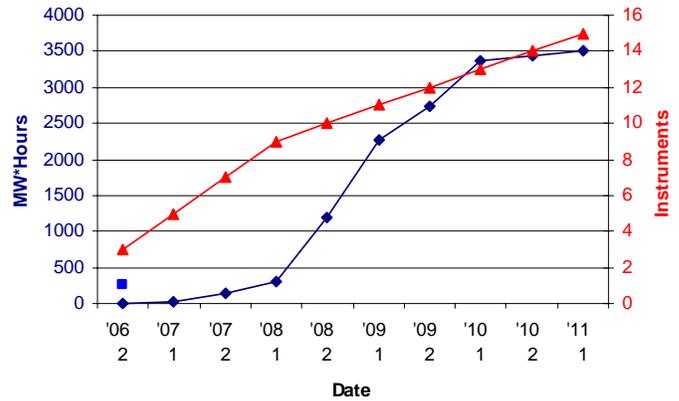


Figure 3: SNS operational ramp-up

### 3. RELATED WORK

Science gateways are evolving into an elegant mechanism with which communities of researchers can be seamlessly integrated with infrastructure, middleware and presentation tools to perform domain specific science [1, 10]. Its physiology includes the following. First is presenting a science community with *application services* such as: presenting a collection of domain specific applications; mechanisms to compose workflows based on them; inducting new applications; visualizing and executing applications in a secure fashion. Second is offering *data services*. This involves the ability to access collections of experimental data held at the facility, download/annotate/share them, and maintain personalized workspaces. Third is providing *infrastructure services*. The gateway should provide transparent access to a computational grid backend, the associated services such as schedulers, directory services, etc. Fourth is enabling access to these services through the use of *security services*. Security services broadly encompass authorization, authentication, auditing and accounting. Finally, a gateway can also include *operational modalities*. For instance, the support of different types of interactions (Web-based, programmatic), disconnected use, etc.

Many domain specific gateways exist, providing selected sets of functionality described above. For instance, the NanoHub [10] provides an application-centric gateway to nanotechnology tools and educational resources to students and practitioners. The Earth System Grid, [11] on the other hand, provides a data-centric view to a community of climatologists enabling them to share and download massive datasets in a reliable fashion. In a similar vein, several other gateway efforts provide community specific accesses to data, applications and computing backend. These include the NEESgrid [12] for earthquake engineers, the LEAD portal [13] for meteorologists and the BioScience gateway [14] projects.

Commonalities exist for many of these gateway efforts in how they compose and present the various services. For instance, parallels can be drawn from the way in which they access the computational

grid backend or data services, etc. Exploiting these similarities, several portal construction toolkits and frameworks have been developed that abstract out routinely performed, generic operations such as job submissions, authentication, data movement, presentation styles, etc. The Open Grid Computing Environment (OGCE) [15] and the Grid Portal Toolkit [16] are examples that enable developers to create grid portals from libraries that support baseline grid technologies. This enables the developer to concentrate on the specialized needs of a particular scientific community. These toolkits themselves tend to adhere to the Java Portlet Specification Request, JSR-168 [17], which lays the foundation for a new open standard for Web portal development. Portlets define an API for building atomic, composable visual interfaces to Web content and services. Implementations of this standard exist in Jakarta Jetspeed [18], IBM WebSphere [19], Oracle Portal [20], BEA WebLogic Portal [21] and GridSphere [22]. While these technologies are available for use, domain users' requirements in terms of graphical interfaces, services, their composition, how they intend to use them should dictate the choice of framework or toolkit. To this end, we conducted a survey, which we present in the next section.

#### **4. SNS REQUIREMENTS ANALYSIS**

We gathered user requirements in three principle ways. First, we conducted several interviews with user groups from SNS and various other neutron facilities to obtain perspectives the essentials for a software solution. Second, we conducted one-on-one demonstrations of prototype software with neutron scientists and gathered feedback on user interface and functionality. Third, several neutron facilities across the world have come together to form the Neutron Science Software Initiative (NeSSI). NeSSI is comprised of several working groups in areas such as data management, data acquisitions, analysis/reduction/visualizations, gateways, etc. We have significant representation and in fact lead or co-lead many of the working groups and use it as a forum to discuss use cases, requirements and share software solutions. The NeSSI has already held three workshops and dozens of tele-meetings in the past two years to discuss and formalize software requirements. Another result of the NeSSI effort is a community-wide general acknowledgement of the necessity of using modern software engineering and the utility of large-scale cyberinfrastructure and high performance computing. Below we list some key specific SNS user requirements derived from these interactions.

Neutron scientists and users fall into three categories based on their levels of expertise. First are novice users who prefer prepackaged feature sets that can be used in a simple, fast, and reliable "push button" fashion. We suspect this is the predominant NSTG use case. Second are seasoned users, who wish to construct and tweak workflows for greater control of their experiments and analyses. Third, is the expert user who is only satisfied with a "shell" so she can compile her code. While users request availability of all of the above, we believe that an elegant and easy to use gateway-based solution will cover in excess of 80% of the use cases. Users desire rich GUI constructs with common look and feel and ease of navigation. The facility will need to provide seamless access to the following: prepackaged, validated neutron science codes; visualization tools; new tools from SNS instrument scientists; batch processing; interactive GUI-based legacy applications; platform dependant applications; web services enabled applications; and stand-alone applications. Thus, we have a variety of tools and applications to support in a software repository, while also exploiting the HPC/Grid availability to our gateway.

Other users are more comfortable running tools on their desktops using their favorite applications. They desire access to raw datasets held at the facility. While SNS cannot seamlessly import their results without validation, we are required to provide programmatic access to these clients. We also need to provide user workspaces to store user local results that are not validated and included in facility endorsed repositories.

User groups require that the facility maintain and archive their experiment and analysis data. This would enable the facility to enforce strong annotation semantics for sharing. The data management system should be able to archive data, provide fast access, maintain replicas for reliability, etc. The SNS raw data rates will be well beyond those handled by any other neutron science facility in the world. Thus, there is the need to efficiently maintain and catalog them for searches.

The SNS data analysis group has conducted extensive interviews with SNS instrument scientists who are responsible for beam line design and construction. The results of those interviews for 13 instruments are summarized in tables that detail the current estimated data rates by instrument. These estimates will be subject to revision as the SNS approaches completion and as the user program ramps up. Table 1 shows the SNS instrument file sizes and rates.

Beam Port	Data Size start-finish (Mbytes)	# File/ Measurement	Time/File	Reduced Data Size (bytes)	User Turnover rate (days)	Min-max production rate (Mb/s)	data rate	visualization transfer rate (Mb/s)
1B	300	1	1h	40k	1-3	0.67		240
2	7 – 13	1	15m	40k	1-3	0.06 – 0.12		5
3	810	1p 10sc	1h	100k	2-3	1.80		648
4A	80	5-10	1s-2h	4x100k	2	0.09 – 643.8		64
4B	80	5-10	1s-2h	100k	2	0.09 – 643.8		64
5	128 – 600	1	10m	10k	3	1.71 – 8.00		102
6	100 – 200	1	1m-1h	1k	2	0.22 – 26.67		80
7	30 – 118	3-4	1s-10m	1k	2-3	0.39 – 946.18		24
11A	148 – 1478	1	5m	1M	2	3.94 – 39.42		118
12	1035 – 2430	8	1h	100k	3-4	2.30 – 5.4		828
14B	384	1	4h	307M	7	0.21		307
17	1373 – 1728	1p 5sc	20m-7d	5.5G	4-5	0.02 – 11.52		1,098
18	900	1p 5sc	20m	3.6G	4	6.00		720
Sum						17.5 - 2,334		4,298

**Table 1: SNS Instrument File Size and Rates.** The “Data Size” column is the size of the raw data assuming all counts are binned (no event mode or sparse arrays). The smaller number is for initial instrument commissioning and the larger number is for later for instance when more detector banks will be installed. The “# File/Measurement” column denotes the number of raw data files that will typically be combined to produce a single reduced file. The “Time/File” is the rate that data files will be produced by the instrument with enough statistics for reasonable analysis. The “User Turnover rate” column denotes an approximate number of days for which a typical user will have the instrument. These requirements can be translated into raw network data transfer rates. These are the range of rates for expected steady flows of data from the instrument to the data management ingestion services. The minimum rate is smallest raw data size divided by the longest measurement time. The maximum rate is the largest raw data size divided by the smallest measurement time. The transfer rate for visualization is the rate necessary to transfer the data to a visualization computer in 10 seconds.

There is a current discussion about whether to store data in event or histogram mode. Table 2 shows a comparison of these two options by instrument. In principle, event mode retains more information and is not subject to sampling bias from binning. However, data rates for the SANS (port # 6 in Table 2) and

Vulcan (port #7) instruments are extreme. Currently, the expectation is that raw data will be stored in event or histogram based on an instrument by instrument Mode optimization.

Port	n/s on sample	n/s measured	bit/s	Event MB/day	histogram MB/day
1B	1E7	7.96E5	5.10E7	535,032	7,000
2	1E7	1.85E2	1.19E4	126	606
3	1E7	2.55E5	1.63E7	176,102	13,608
4A	1E9	1.13E4	7.21E5	7,762	4,024
4B	1E9	1.13E4	7.21E5	7,762	57,263
5	1E7	3.82E5	2.45E7	256,815	17,184
6	1E9	6.98E7	4.47E9	48,230,400	5,000
7	2E8	4.78E6	3.06E8	3,256,051	547,651
11A	3E7	8.36E5	5.35E7	577,834	42,578
12	1E7	1.59E5	1.02E7	110,064	24,840
14B	5E7	6.41E4	4.10E6	44,301	2,304
17	1E6	9.55E3	6.11E5	6,384	73,498
18	3E6	7.41E4	4.74E6	49,110	47,769

Figure 4a gives the data production rate of SNS in terms of GB/day as a function of year. This is calculated using the information in Table 1 assuming that there are no pauses between each file being produced. The increase in rate is attributable to instruments coming online in each year. The majority of reduced data can be attributed with beam lines 17 and 18. The green “old raw” curve is a previous estimate from 24 months ago. It is included to demonstrate how the facility’s data estimates have changed over time. Figure 4b plots the total, cumulative data storage needed to keep one copy of the raw data and one copy of the reduced data. Figure 4c gives

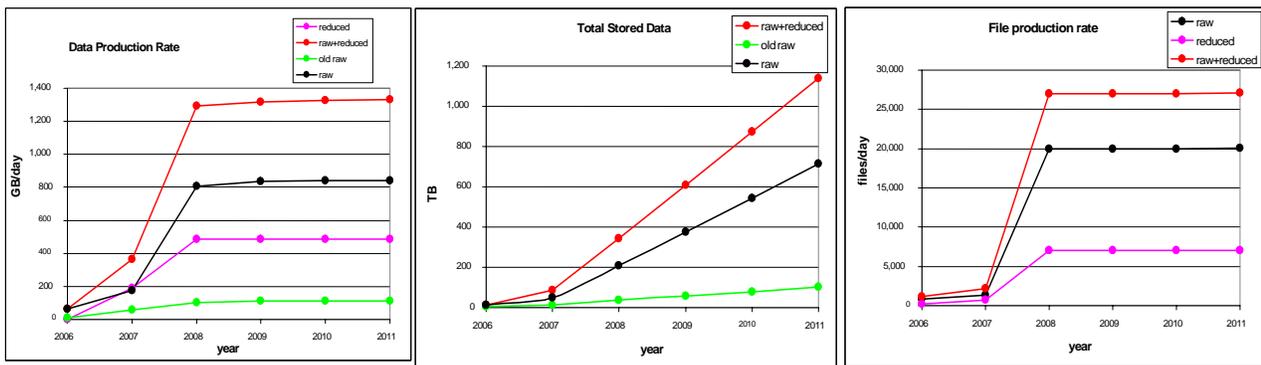
**Table 2: Data Rate Comparison Event/Histogram Mode**

the number of files produced per day as a function of year. This is to characterize how much information will be added into a database cataloging experiments or used to locate data.

These data requirements for SNS are often succinctly summarized by the following figures of merit:

- By, 2008, data will be generated at the rate of on the order of a TeraByte per day contained in 20,000 to 25,000 files
- By 2011, the cumulative data store is expected to be 1.2 PetaBytes.

These requirements are a significant input for planning for the NSTG and the SNS data analysis infrastructure. We also believe that as other next generation sources proceed with construction, they will find similar data requirements for their new beam lines.

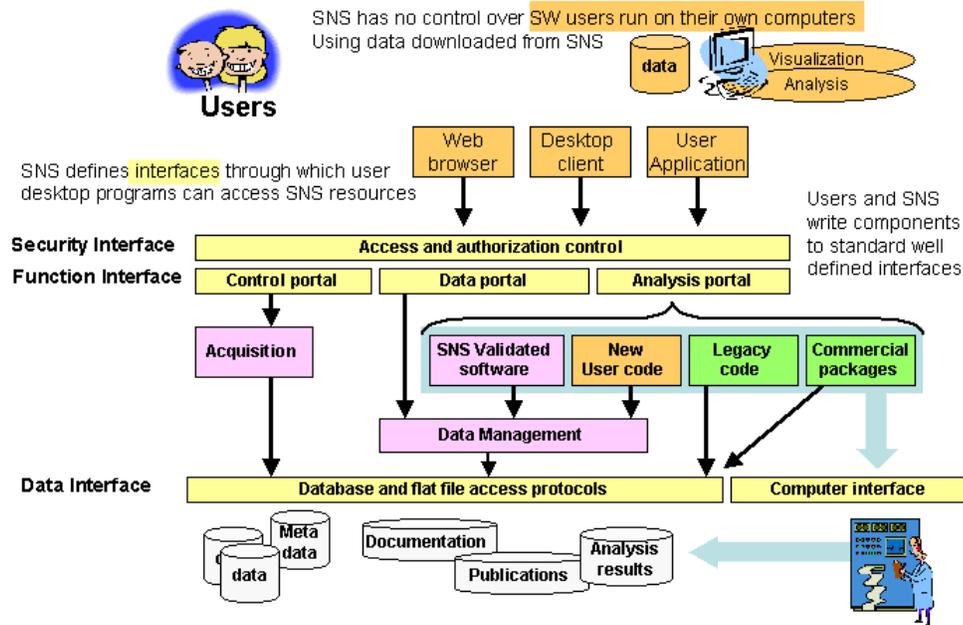


**Figure 4: a,b,c SNS data production rate estimates**

## 5. SNS SOFTWARE SERVICES ARCHITECTURE

Based on requirements expressed by users, our solution needs to enable two primary modes of operation. First, the user accesses SNS services through a browser connected to a portal infrastructure

that presents a compilation of backend software services. Through this gateway, neutron scientists can perform data functions of acquisition monitoring, reduction, analysis, pipelining, visualization and sharing with collaborators. Second, the user requires programmatic access to the same services so she can use her favorite desktop client tool or legacy application. Figure 5 depicts the software interfaces comprising of several hierarchical layers capturing the desired use cases. In each case, users pass through a security interface comprised of authentication/authorization services to negotiate access and then perform operations on resources (data, computation, database, instrument, etc.) through corresponding function-specific interfaces.



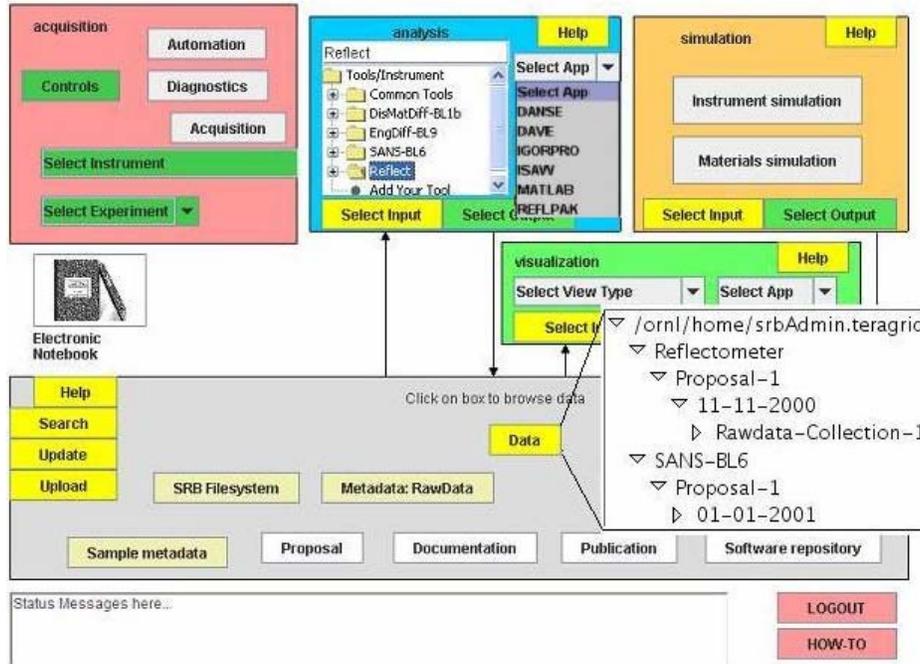
**Figure 5: SNS Functional Interfaces Architecture**

Below, we describe the neutron science portal, its design and architecture based on three key functional interfaces namely: Data Management, Application Management and Security. We describe the internal architecture of the aforementioned based on a prototype portal implementation.

**The NSTG:** The prototype NSTG portal (Figure 6) is implemented as a JAVA [23] applet tying together backend SNS services such as data/metadata management, visualization/simulation capabilities and application management. The portal server is built atop the JAVA-based Tomcat server engine [24], version 5, and the Jetspeed portal framework [25], version 1.1.3. Jetspeed portal framework is primarily used for servlet/jsp executions, user customizations and state management. Web services support is implemented through the use of Apache Axis [26] and JAX-RPC [27]. A neutron scientist is only required to have a JAVA Runtime Environment (JRE) [23] enabled browser to access the portal and its services. The prototype portal brings with it several sophisticated features such as:

- Ability to construct and launch simple pipelines of analysis tools for batch processing
- Downloading of software tools, adding software tools from user's desktop into a pipeline that is being constructed on the portal, etc.
- Execution of third party, interactive applications such as a MATLAB [7], ISAW [8], etc., from the Web portal
- Browsing and downloading of datasets maintained in a Storage Resource Broker (SRB) [28, 29] collection (which in turn maintains data in HPSS [30] archives, all transparent to users.)

- Launching of simple visualizations and static image views by performing 1-D, 2-D, 3-D and histogram plots delivered through backend JAVA server page technologies [31]
- Per-user electronic notebooks to annotate analyses
- Searching and updating metadata tables through database queries
- Job and data operations can be performed under group as well as individual accounts and several other features.



**Figure 6: The Neutron Science TeraGrid Gateway showing analysis tools, SRB data browsing visualization and metadata features**

At the backend, the NSTG runs on the ORNL-TeraGrid Linux cluster with SRB access to data held in HPSS and NFS file systems, metadata in an ORACLE database with JDBC support and several neutron science specific pre-installed binaries and software packages. It further has access to grid resources offered by the TeraGrid at large.

**Data Management: (yellow boxes)** We have designed and implemented the data management system as a hierarchical layered set of services as shown in Figure 7. The data management subsystem of SNS (denoted as yellow boxes in Figure 7) provides a service oriented architecture for its clients. It comprises of several components like the file creation service, file read service, metadata management/query service, databases, the SRB and several administration services. Databases include metadata regarding credentials, users, raw datasets, proposals, experiments, etc. Raw datasets and user workspace data are maintained by SRB as separate collections. Each of the data management components is implemented as SOAP web services [32]. Web services offer a highly inter-operable environment and support a wide variety of data types including SOAP objects and XML data.

The data management system has a set of higher-level clients whose needs are known such as the Application Manager (purple boxes) or Data Acquisition Services (orange/brown boxes). To these known clients, it provides a set of well-defined interfaces and libraries for invoking its web services.

The benefits are as follows. First, these higher-level components can remain oblivious to the construction of SOAP messages while interacting with web services since the library handles them. Second by providing a specific client library for each higher-level component, we ensure that the application code is independent of the data management system implementation. When a web service behavior changes, the client library is changed to retain the functionality without modifying its interface. Thus the user application need not be modified or recompiled as long as the functionality is still supported. All other clients interact with the data management subsystem through its well defined web services. In addition, we have implemented a service API for SRB that exports the functionality of SRB as stateful web services. The web services are implemented using Apache Axis and JAX-RPC. Uploads and downloads are implemented using HTTP servlets. The web services and servlets interact with SRB using 'Jargon' - the JAVA API for SRB.

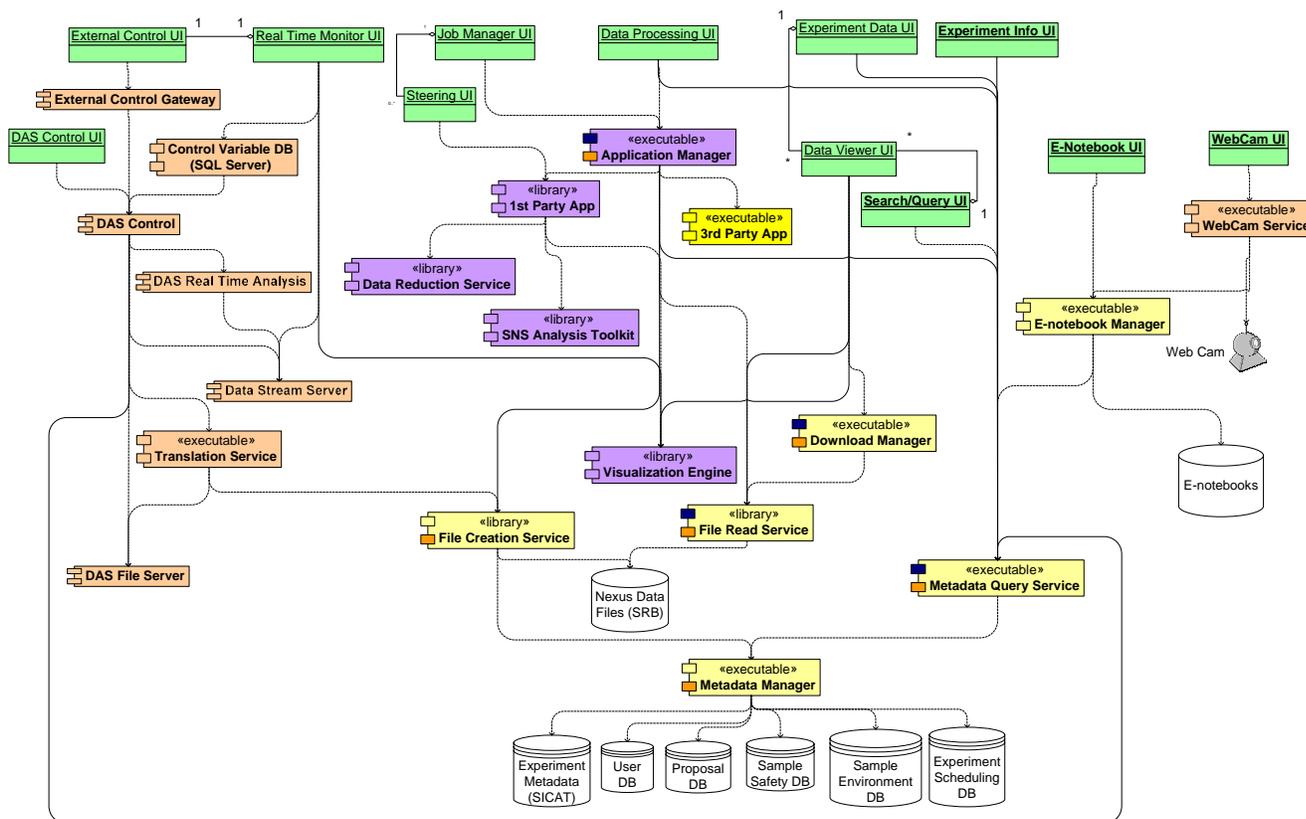


Figure 7: Overall Software Architecture

**Application Management: (purple boxes)** The Application Manager is responsible for orchestrating application executions. Figure 7 shows the position of the application manager in the overall architecture. Applications include both 1st and 3rd party packages and can be executed either interactively or in batch mode. To execute applications, users interact with the user interface (UI) (Figure 7) or programmatically invoke UI services to specify parameters, input/output files and tools that are then parsed into an executable command-line or script. The application manager internally, determines required and available resources and stages executables and data appropriately for lower level applications. The application manager orchestration also includes the creation of static or dynamic workflows and compositions. This creates a simpler environment in which the individual applications can be created in an isolated environment. We expect that the neutron science community

will create many of the applications (or import them from existing third party applications). To the extent that the application manager can simplify application development, it will accelerate the inclusion of advanced discipline specific application functionality into the overall architecture. In the current prototype portal implementation, jobs are executed on the local TeraGrid cluster using local resource management tool such as, PBS, after obtaining data from the data management service. Jobs submissions can be achieved using standard grid job execution techniques (Condor-G [33], Globus-job-run [34], Gridshell [35], etc.). In the future, we plan to extend this mechanism through the use of standard portlets from the OGCE Grid portal framework. This would also enable us to perform intelligent scheduling through the use of grid information services.

Web-based executions of interactive third-party applications such as, MATLAB, [7] follow a similar process but with a slight variation. With such applications there is the need to efficiently forward the X display back to the user's desktop. This is further complicated since an application can be launched anywhere in the grid. To address this, we use a mechanism such as WiredX [36] or VNC [37] that delivers a JAVA-based X window system to the user's browser that elegantly achieves secure X forwarding. Using this approach, we can deliver standard UNIX applications to non-UNIX desktops.

In our prototype implementation, we have demonstrated the execution of tools provided by the facility. One can imagine an advanced users wishing to upload and run their own tools on the facility's computing resources. To address the security ramifications therein, we will explore sandboxing approaches such as executing all applications on a virtual machine.

**Security:** We are given the requirement that anything more than a simple username/password-based authentication is beyond the current comfort zone of our user base. Any user credential and certificate mechanism required by HPC resources and grid environments will need to be transparently maintained and propagated by the authentication/authorization service of the portal. In our prototype, we use GSI proxy credentials for authenticating with the SRB subsystem. The SRB credentials are stored within the authentication service for all users and are transparently used by the services while accessing the SRB system. The SRB server is itself started with the host certificate on a TeraGrid node. We intend to extend the use of proxy credentials to submitting analysis jobs on other TeraGrid nodes.

## 6. DISCUSSION and CONCLUSION

The prototype portal is operational under a "demo" mode and has been available for use by neutron scientists from SNS and other neutron facilities across the world (ISIS in the UK; JPARC in JAPAN; IPNS, Lujan Center, NIST NCNR and HFIR in US). Something we have realized as part of the collaborative development efforts is that the currency of collaboration is documentation, not code sharing or other means. Reports from initial user experience suggest that they welcome the idea of having the facility cater to everyday neutron science needs. The notion of gateways, portals and network-based computing for analyses, in general, is still new and novel within the neutron community, unlike say, high-energy physics or biology, where such ideas are commonplace. This suggests that the onus is on the gateway/portal, then, to make a neutron scientist's experience as easy as possible. First, requiring very little—to zero new software installations—on the user's desktop is an almost absolute requirement. Second, the gateway should be inclusive of users' favorite tools (be it batch or interactive). Users should be able to add new tools (although security policy is an issue). Third, the portal should

address latency issues. Users will simply refrain from using features that they think are slow: say, workflow wiring, etc. What do the aforementioned mean to some software design decisions then?

For instance, from an authentication standpoint, since users will not bear anything more complicated than username/passwords the portal will need to perform any kind of mapping required to access HPC/Grid resources (say building upon MyProxy [15] server like solutions). From an analysis standpoint, we can be agnostic to interactive or batch tools, but interactive applications are latency intolerant. Here, we are exploring compression schemes that come with X forwarding systems. From a visualization perspective, we can attempt many intelligent data caching policies or perform parts of the rendering on user desktops, while still maintaining the balance between latency and client-side installations. JAVA Webstart like solutions do alleviate the problem a little, but are language specific and tend to violate the “require no client installation” requirement. Similarly composing GUI-based workflows can sometimes increase latency, but offloading those to client desktops can buy flexibility allowing users to work on them at leisure and upload the pipeline when finished. These arguments suggest that there is the need for both lightweight and rich GUI interfaces and the portal should handle them intelligently.

As users begin to use the portal extensively, they are likely to add new tools and applications. Automatically adding and configuring third party applications to portals has recently been studied at length in grid communities. We can benefit from their progress. We are currently working on several of these issues and plan to reuse and build upon the wealth of development that is underway in portlets (JSR 168 [17] compliant tools that can be reused). Another point of constant debate is that of sharing analysis results and maintaining pedigree information and workflow metadata. With the gateway-model of performing neutron science analyses, users are a lot closer to sharing data than they otherwise would have been.

In summary we are at a requirements definition phase. These requirements exercises have been informed by the experience of extensive development and demonstration of a prototype portal infrastructure. The science user community has welcomed this approach and believes that it holds promise. Accordingly, we are concluding the requirements phase and moving rapidly to specification and implementation in anticipation of production deployment.

## **ACKNOWLEDGEMENTS**

We would like to acknowledge additional useful conversations and contributions from the NeSSI collaborative group including R. McGreevy, D. Mikkelson, R. Mikkelson, T. Worlton, as well as W. Heller and the SNS instrument scientists. We would also like to acknowledge the research sponsorship for these activities. Research sponsored by the U.S. National Science Foundation under interagency agreement DOE No. 0700-S664-A1, NSF Cooperative Agreement ACI-0352164 and Cooperative support agreement No. ACI-0338605 and by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory (ORNL), managed by UT-Battelle, LLC for the U. S. Department of Energy under Contract No. DE-AC05-00OR22725. Research executed under U.S. Department of Energy Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC. The submitted manuscript has been authored by a contractor of the U.S. Government under Contract No. DE-AC05-00OR22725. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish

or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

## 7. REFERENCES

- [1] TeraGrid. <http://www.teragrid.org/> [15 August 2005].
- [2] Spallation Neutron Source. <http://www.sns.gov/> [15 August 2005].
- [3] Mason T. E, Taylor A. D. Neutron scattering in materials research. *MRS Bulletin* December 1999; 24(12):14.
- [4] J-PARC. <http://j-parc.jp/MatLife/en/index.html> [15 August 2005].
- [5] ISIS. <http://www.isis.ac.uk/TargetStation2/> [15 August 2005].
- [6] Catlett C. *The TeraGrid: A Primer*. The TeraGrid Project: <http://www.teragrid.org/about/TeraGrid-Primer-Sept-02.pdf>, September 2002.
- [7] MATLAB. <http://www.mathworks.com/products/matlab/> [15 August 2005].
- [8] DAVE. <http://www.ncnr.nist.gov/dave/> [15 August 2005].
- [9] ISAW. <http://www.pns.anl.gov/computing/isaw/> [15 August 2005].
- [10] Goasguen S, Madhavan K, McLennan M, Lundstorm M, Klimeck G. The NanoHUB: A Science Gateway for the Computational Nanotechnology Community. *Workshop on Science Gateways: Common Community Interfaces to Grid Resources, GGF-14*, June 2005.
- [11] Earth System Grid. <https://www.earthsystemgrid.org/> [15 August 2005].
- [12] NEESGrid. <http://www.neesgrid.org/index.php> [15 August 2005].
- [13] Droegemeier K, Chandrasekar V, Clark R, Gannon D, Graves S, Joseph E, Ramamurthy M, Wilhelmson R, Brewster K, Domenico B, Leyton T, Morris V, Murray D, Plale B, Ramachandran R, Reed D, Rushing J, Weber D, Wilson A, Xue M, Yalda S. Linked environments for atmospheric discovery (LEAD): A cyberinfrastructure for mesoscale meteorology research and education. *20th Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, January 2004.
- [14] Blatecky A, Gamiel K, Ramakrishnan L, Reed D, Reed M. Building the bioscience gateway. *Workshop on Science Gateways: Common Community Interfaces to Grid Resources, GGF-14*, June 2005.
- [15] OGCE. <http://www.ogce.org/index.php> [15 August 2005].
- [16] Dahan M, Boisseau J. R. The GridPort toolkit: A system for building grid portals. In *Proceedings of the 10th IEEE international Symposium on High Performance Distributed Computing (HPDC)*, 2001. IEEE Computer Society: Washington, DC, USA.
- [17] Sun Microsystems, JSR 168 - The Portlet Specification, Java Specification Requests, October 2003. <http://www.jcp.org/en/jsr/detail?id=168>
- [18] Jetspeed. <http://portals.apache.org/jetspeed-2/> [15 August 2005].
- [19] IBM Websphere. <http://www-306.ibm.com/software/websphere/> [15 August 2005].
- [20] Oracle Portals. <http://www.oracle.com/technology/products/ias/portal/index.html> [15 August 2005].
- [21] BEA Weblogic Portal. <http://www.bea.com/content/products/weblogic/portal/index.htm> [15 August 2005].
- [22] GridSphere. <http://www.gridsphere.org/> [15 August 2005].
- [23] Java & Java Runtime Environment (JRE). <http://java.sun.com> [15 August 2005].
- [24] Tomcat server engine. <http://jakarta.apache.org/tomcat/> [15 August 2005].
- [25] Jetspeed portal framework. <http://portals.apache.org/jetspeed-2/> [15 August 2005].

- [26] WebServices- Axis. <http://ws.apache.org/axis/> [15 August 2005].
- [27] Java API for XML-Based RPC (JAX-RPC). <http://java.sun.com/xml/jaxrpc/index.jsp> [15 August 2005].
- [28] Baru C, Moore R, Rajasekar A, Wan M. The SDSC Storage resource broker. *Proceedings of the CASCON'98 Conference*, November 1998. and SRB website <http://www.sdsc.edu/srb>
- [29] Rajasekar A, Wan M, Moore R. MySRB & SRB: Components of a data grid. *The 11th International Symposium on High Performance Distributed Computing (HPDC)*, July 2002. IEEE Computer Society: Washington, DC, USA.
- [30] Watson R. W, Coyne R. A. The parallel I/O architecture of the high-performance storage system (HPSS). In *Proceedings of the 14th IEEE Symposium on Mass Storage Systems*, September 1995. IEEE Computer Society: Washington, DC, USA.
- [31] JavaServer pages Technology. <http://java.sun.com/products/jsp/> [15 August 2005].
- [32] Snell J, Tidwell D, Kulchenko P. *Programming Web Services with SOAP* (1st Edition), O'Reilly: 2001
- [33] Frey J, Tannenbaum T, Foster I, Livny M, Tuecke S. Condor-G: A Computation Management Agent for Multi-Institutional Grids. *Journal of Cluster Computing* volume 5. pages 237-246. 2002.
- [34] Czajkowski K, Foster I, Karonis N, Kesselman C, Martin S, Smith W, Tuecke S. A Resource Management Architecture for Metacomputing Systems. *Proc. IPPS/SPDP '98 Workshop on Job Scheduling Strategies for Parallel Processing*. pg. 62-82. 1998.
- [35] Walker E, Minyard T, Boisseau J. GridShell: A Login Shell for Orchestrating and Coordinating Applications in a Grid Enabled Environment. *Proceedings of the International Conference on Computing, Communications and Control Technologies*. pp. 182-187. Austin, Texas. August 2004.
- [36] WiredX. <http://wiredx.net/> [August 15, 2005]
- [37] VNC. <http://www.realvnc.com/> [August 15, 2005]