

Building a Data Grid for the Australian Nanostructural Analysis Network

Brendan Mauger, Jane Hunter, John Drennan, Ashley Wright, T. O'Hagan
The University of Queensland,
St Lucia, Queensland, Australia
jane@itee.uq.edu.au

Abstract:

This paper describes the architecture and services developed by the GRANI project for the Australian Nanostructural Analysis Network Organization (NANO). The aim of GRANI was to provide the NANO community with a scalable, distributed data management solution and a secure collaborative environment to ensure high speed access to and seamless sharing of their data, instruments, analytical services and expertise. A grid-enabled, distributed system was developed that links the major Australian microscopy instruments to an underlying distributed national imagery database, a network of microscopy experts and image processing and analytical services through an authenticated Web/Grid Portal. The aspects that are particularly innovative and that will be described in depth include:

- *The Nano Image Database (NIDB) –an indexed, distributed archive of images captured directly from the advanced instruments and copied to the National Data Facility using Gridftp;*
- *Combining the Australian Partnership for Advanced Computing (APAC)'s high performance computing (HPC) facilities and grid environment with the Kepler workflow system to enable high speed file movement, image analysis and 3D reconstruction;*
- *Real-time video conferencing and video annotation services to improve support for remote access to advanced microscopy instruments and experts.*

1. Introduction

Scientists from across the biological, materials and chemical sciences are increasingly employing advanced microscopy and characterization techniques to help them understand the nanostructure of inorganic and organic materials in order to solve complex biomedical, scientific and engineering problems. In the process, they are generating massive volumes of multi-disciplinary image (both 2D and 3D) and video data. Advances in microscopy techniques such as atom probes and 3D cryo-electron microscopes, have increased the speed, resolution, dimensions and scale at which images are being generated. Scientists and microscopy centres are struggling with the problem of efficiently managing, processing, sharing, indexing and retrieving these large image collections generated by distributed virtual teams.

The aim of the ARC-funded GRANI (Grid-enabled Archive of Nanostructural Imagery) project [1] was to provide the Australian Nanostructural Analysis Network Organization (NANO) with a scalable, federated, distributed data management solution – a secure Web portal to a Grid-based image archival and analysis system. Central to this solution is the NANO Image Database (NIDB) - a large-scale, distributed data management system. Images and associated metadata are captured directly from the instruments. Users can then selectively upload and store required images in a local image server node of the NIDB. Long term archival is supported by storing a copy of the images at the National Data Facility. A Web interface provides searchable access to the images and also allows users to define access privileges and Creative Commons Licenses [24] for specific users and images. A comprehensive set of metadata is captured with the images, enabling advanced search and intelligent image matching. Figure 1 illustrates the web-enabled portal which provides access to stored data/images, expertise, instruments, analytical tools and annotation services.

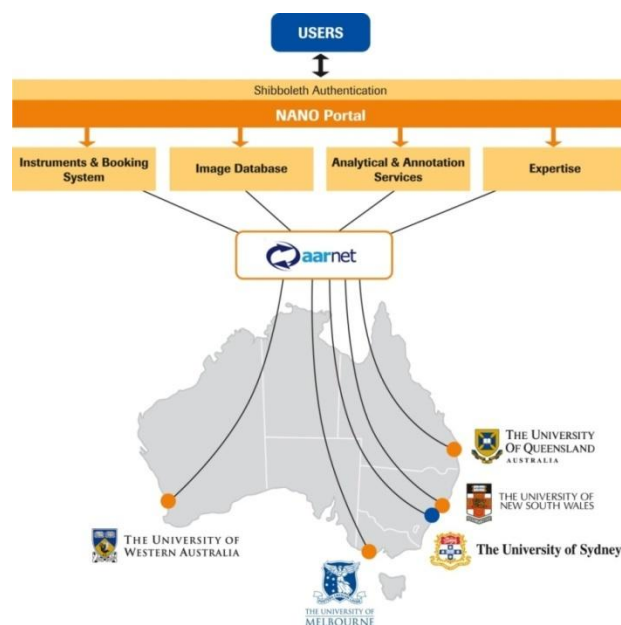


Figure 1 : Web Portal for the NANO Network

The remainder of the paper is structured as follows: Section 2 describes previous work and related projects; Section 3 describes the overall system architecture and technologies employed; Section 4 describes the user interface; Section 5 describes the tele-microscopy services; Section 6 outlines future work plans and conclusions.

2. Related Work

A number of projects have been developing Grid-based tools and infrastructure to assist research communities with managing and analysing large volumes of images captured from advanced scientific instruments. The Large Hadron Collider (LHC) being built at CERN near Geneva has extensively employed a grid based approach, LCG [2], to provide scalable infrastructure for international collaborators. MEDIGRID [3] is a French project exploring the use of the GRID technologies for tackling the processing of huge medical image databases. The NeuroGrid [4] project aimed to provide image archival, curation and analysis capabilities for the neuroimaging community using Grid technology. . GridPACS - a Grid-based Image Archival and Analysis System [6,7] is an example of a distributed XML-based image management system. SIDB (Scientific Image Database) – is open source software for archiving 2D and 3D microscopy images [8]. The BIRN (Biomedical Informatics Research Network) [9] project focuses on collaborative access to and analysis of images and datasets generated from neuroimaging studies. It uses the Storage Resource Broker [10] for the distributed data management middleware layer. eDiamond [5] targets deployment of Grid infrastructure to manage, share, and analyze annotated mammograms captured and stored at multiple sites. The Open Microscopy Environment (OME) [11] produces open tools and adoptable XML based standards to support data management for biological light microscopy.

The GRANI project has surveyed and evaluated these previous related projects and adopted and integrated those components that satisfy the requirements of the NANO community (documented through a comprehensive user requirements survey) and that we believe will provide the most robust, extensible and scalable framework. In particular we adopted OME for the core metadata schema, a relational database (MySQL) for the metadata store which contains URIs to multiple copies of the files. At this stage, we have chosen not to use SRB because of concerns related to its scalability, stability and robustness. We use the Grid (GridFtp) to transfer replicated files to the location of data processing and to the National Data Facility (at the Australian National University in Canberra) for long term archival. Distributed computing facilities are

available via the APAC Grid for the high-speed image processing. In addition, we have extended remote or tele- microscopy work that began at the University of Queensland in 2000 [12] as an outreach program to secondary schools in remote regions. We have refined and extended this work to provide a real-time annotation service for high-resolution video streams and a backend image database to enable storage of high res images captured during tele-microscopy sessions – these extensions support the more sophisticated requirements of the multidisciplinary research communities who use the NANO Major National Research Facility (MNRF).

3. System Architecture

Figure 2 illustrates the overall architecture and technologies for a single node of the NANO Distributed Image Data Store. Figure 3 illustrates the national distribution of networked characterization laboratories and instruments, regional storage nodes and the National Data Facility (NDF) in Canberra. Files originate from a particular instrument in a Lab and are then uploaded to the local node of the NIDB using the secure web interface (Shibboleth). A MySQL database is used to store the metadata, file indices and information that is needed by FAMS (File Access and Management System). FAMS provides the interface between the NIDB node and the Grid environment and manages file movement across the Grid using Gridftp [25] and RFT (Reliable File Transfer). Instrument-specific post-processing workflows are applied to the captured files to extract metadata. These are designed so that they can easily be customized to support new instruments or perform additional compute-intensive tasks (such as segmentation) using grid HPC facilities.

3.1 PHP Interface & Web Technologies

The dynamic web site has been programmed using PHP. PHP was chosen because it is platform independent, Web-centric [13] and a well-supported language enabling rapid deployment across the diversity of environments and platforms that exist within the NANO community. The Web Portal provides a single federated user interface to the locally deployed web sites and storage nodes. Security is provided via Shibboleth user identification and authorization which authenticates users across institutions via the institutional identity providers. Thumbnails are dynamically generated from the captured high-res files using Image Magick [14]. PECL PHP libraries provide SSH access to the live lab work areas and file systems. AJAX, XML, Javascript and CSS are the underlying technologies that ensure dynamic web interfaces, scalable high speed performance and highly responsive interactivity.

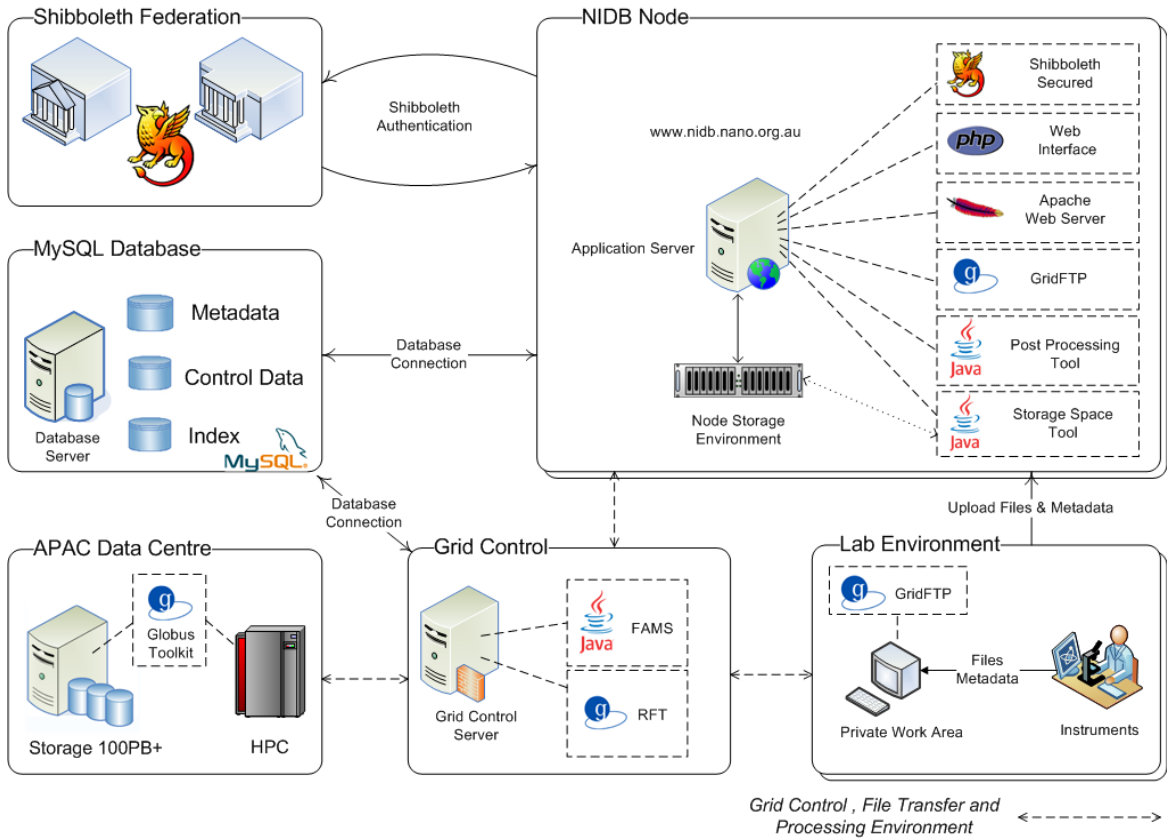


Figure 2: System Architecture

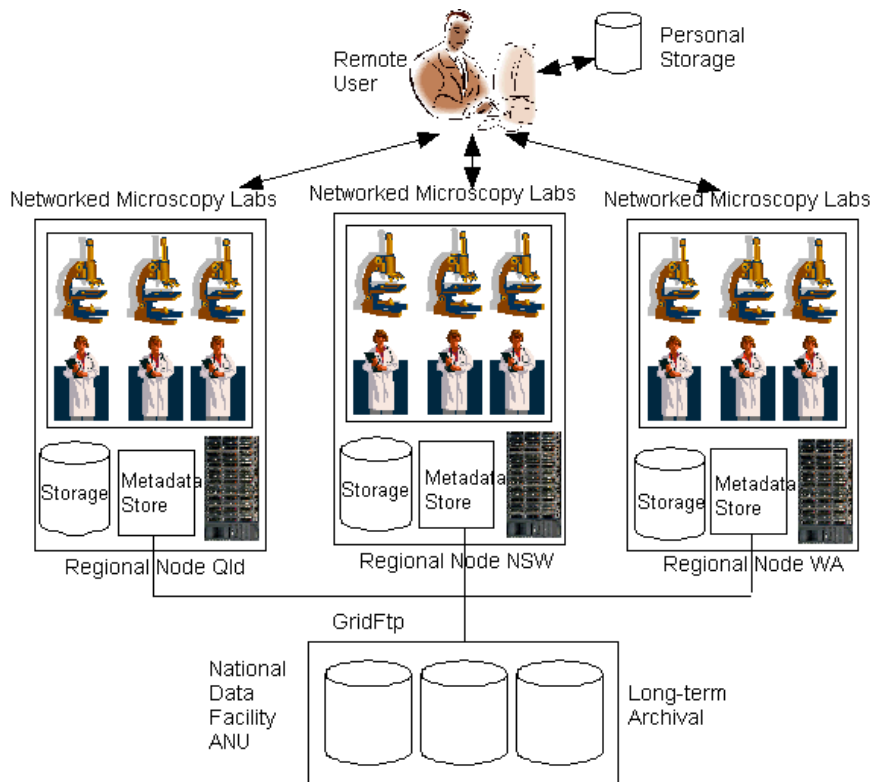


Figure 3: National Overview of Architectural components

3.2 Relational Database and Metadata

A MySQL database is used for storing the structured (XML) metadata descriptions. An extensible XML schema (based on OME) was designed to document the metadata captured from a wide variety of instruments. There are three levels of metadata: generic, instrument-specific and extensions. Generic metadata comprises those attributes that all image files possess e.g., *file name, title, creator, date, instrument, project, sampleId, discipline, topic*. Instrument-specific metadata (e.g., “Scanning Electron Microscope Metadata”) contains attributes that are common among a class of instruments e.g. *scan speed, Micron Marker, magnification, working distance, accelerating voltage, spot size*, etc. The last class of metadata, “Extensions”, contains all the remaining metadata generated by the instrument plus any ad-hoc project-specific metadata. The generic and instrument-specific metadata values are stored in structured, indexed tables, ensuring fast search and retrieval. The metadata extensions table enables additional flexibility and adaptability in the metadata schema. The “File Index” field maintains a record of files and their locations. The “Control Data” field maintains a record of files that need to be moved or processed by the grid environment. At this stage we have chosen not to use the Storage Resource Broker (SRB) for storing and managing files due to concerns regarding the performance of MCAT and issues related to SRB’s robustness and user interface.

3.3 GridFTP and the FAMS

The Nano Image database makes use of a number of grid technologies to manage the transfer of files between NIDB Nodes and the APAC National Data Facility. A GridFTP server, a parallel version of FTP for the Grid [25], is installed on every node and data store and is used for transferring files between locations. The Globus 4 Reliable File Transfer service is used for managing the numerous GridFTP sessions, and ensuring that failed transfers are restarted. The RFT service runs on an APAC Gateway machine [16] and is independent from any of the nodes. There are numerous RFT services which could be used, providing a degree of tolerance against machine failures. Using the combination of GridFTP and RFT provides a high performance and reliable backbone for transferring data between nodes and institutions around Australia.

In order to connect the NIDB web portal with the RFT service, an extensible Java tool, known as FAMS, has been written. FAMS connects to the MySQL database and reads control information which has been stored there by other parts of the NIDB database. FAMS makes use of two different sets of control information, the first is the queue of

files to transfer. This table is either created by users manually moving files via the web interface, or through an automated tool like POSTP (see section 3.4). The second set of control information is the queue of files to delete. FAMS uses this control information to schedule transfers and deletes with the RFT service. Files are scheduled in batches, so that each image is transferred together with any support files. RFT permits a transaction based approach to transferring files - either all files arrive at the destination, or none arrive. FAMS makes use of these transactions to ensure that any accompanying files (see section 4 – *The Upload Interface*) always remain with the image. During a transfer all status information is written to the database. FAMS is run as a background process, continually monitoring the database for new information, as well as monitoring the status of active transfers. Because the tool only performs monitoring, it consumes minimal CPU time, and only modest memory resources.

3.4 Automated Post Processing and Metadata Extraction

A post-processing step streamlines the metadata extraction from the instrument/file header and ensures that the images and metadata are validated through a quality control process. This semi-automated data curation step relieves the user from the tedious process of manually entering large amounts of metadata. The flexible metadata schema that we have developed can be completed by parsing and mapping the file header generated directly from each instrument. An extensible Java tool, known as POSTP uses information stored in the MySQL database to determine which files and data require processing. The Java-based framework also allows for easy integration with grid-based compute and storage facilities. POSTP carries out common processing tasks such as file compression and manipulation, as well as instrument-specific analytical algorithms and metadata parsers. POSTP runs continuously in the background, actively utilizing available resources to perform compute intensive tasks on user’s behalf. An example of a task that is performed on a regular basis would be the extraction of metadata from file-headers and ingestion of the metadata in the database.

3.5 Dynamic Storage Space Management

The scientific community is producing an exponential amount of data each year. Many current storage solutions are inadequate, non-scalable and don’t support long term data storage. Each node has a potentially finite amount of storage space, serving users producing a potentially infinite amount of data. A Java tool has been developed that manages the local storage space available at each node. The

configurable storage space tool monitors the local storage resources to ensure that there is always space available for new data. A defined upper and lower limit is used to trigger and stop file archival. The tool uses access frequency, extracted from the MySQL database to determine which files should be scheduled to be moved to the data centre for archival. The tool also provides a mechanism for deleting inconsistent and obsolete data, such as unused thumbnails.

3.6 Identity Management and Shibboleth

Shibboleth is used to provide federated identity management and authentication of users. Shibboleth allows users to sign-in to the system and access nodes at external institutions, using the credentials provided from their home institution. It enables transparent, seamless secure access to multiple sites. However as Lorch et al [17] explains, Shibboleth does not provide a comprehensive and dynamic solution for a Grid environment. Grid and application resources are made available to the user through a controlled agent-based approach and Grid system certificates (provided by the APAC Certification Authority (CA)). To overcome the problems of poor interoperability between Shibboleth and Grid authentication, we have adopted a workable agent-based approach in which the application is provided with access to all available resources and manages and performs actions on behalf of the user. Current trends are moving towards a more comprehensive integration of grid access and Shibboleth similar to the GridShib project developed by David Spence et al [18] and the Shebangs project [19], which may offer a promising alternative in the future.

4. NIDB User Interface

In a typical workflow, files originate in a “Lab Environment” where the scientists/users conduct their experiments. The system enables the scientist/user to save their files along with the metadata directly from the instrument into their private workspace.

The lab environment: The lab environment provides a small temporary workspace. The user can access view and access their files in any of the lab environments within any of the nodes, all through the single integrated and secure Web portal. After selecting a particular lab environment, the interface illustrated in Figure 4 is displayed. To upload a file to the NIDB, the user clicks on the upload link (arrow) under the NIDB column (Figure 4). The “priority” combo box on the top right can be used to assign a priority which is dynamically managed by the FAMS.

The upload interface: The Image Upload interface, Figure 5, is used to complete an upload request. The interface allows any “Accompanying Files and

Folders” to be uploaded and saved with the primary file. For example, in Figure 5, the file “fract 30KV.txt” (main file) also has two folders and two files associated with it.

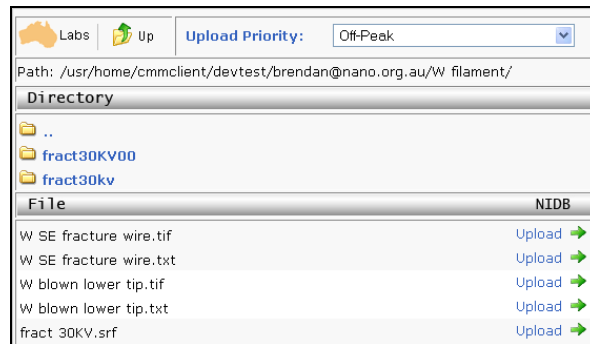


Figure 4: View of Files in a Lab Environment

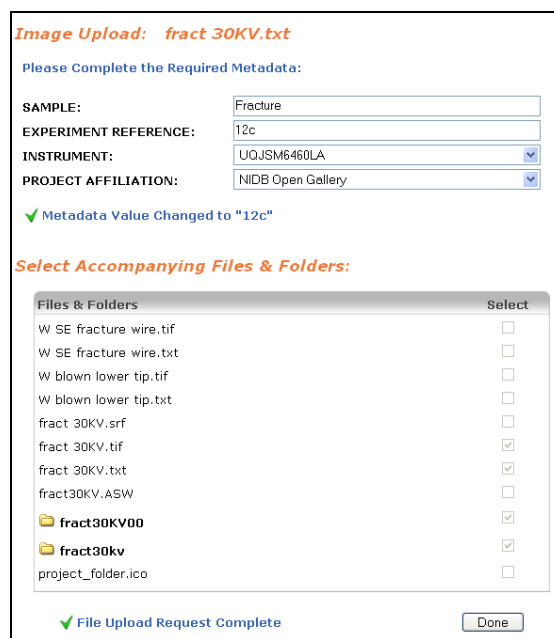


Figure 5: File Upload Interface

The edit interface: The metadata associated with files that have been uploaded to the database can be updated or edited. The editing interface, shown in Figure 6, allows the user to define access privileges, edit metadata as needed and attach Creative Commons Licenses [24] to their work.

The Search Interface: A secure interface has been developed to allow users and their collaborators to search the database via the metadata. The search fields depend on the metadata generated from each instrument e.g., a JEOL SEM (Scanning Electron Microscope) generates different metadata to an FEI TEM (Transmission Electron Microscope). Metadata search fields can be chosen from a drop down list.

Figure 7 illustrates the auto complete feature which expedites metadata input. The list will dynamically update to show only those terms that contain the current input letter sequence.

Only those images that the user is permitted to view (as defined by the access policies attached at file ingest) will be retrieved from the NIDB.

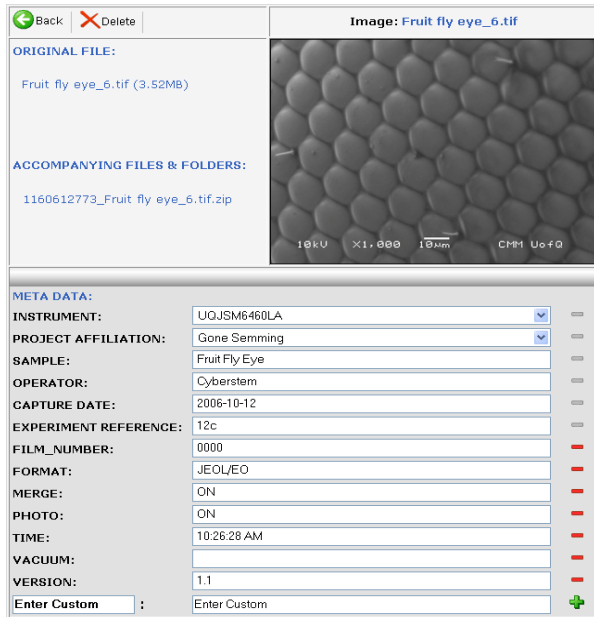


Figure 6: Interface for Editing File Metadata

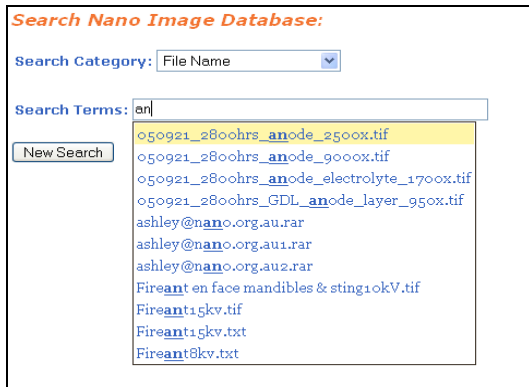


Figure 7 : Search with Auto-Complete

Figure 8 illustrates returned results. The “Search Current Results Only” button can be used to refine searches. This feature will only search files that have been returned by the previous search, allowing users to build up comprehensive and specific searches.

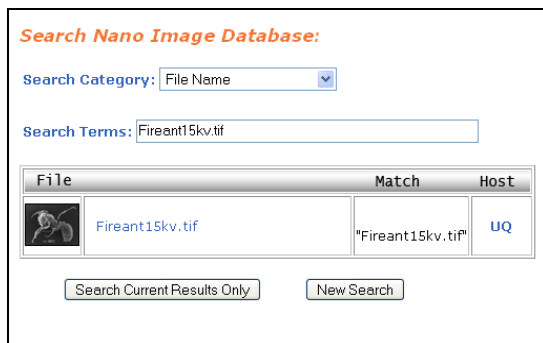


Figure 8: Positive Search Match

Viewing Files: Figure 9, is an example of the interface provided for viewing images - the example image is a fractured tungsten cathode. This interface presents a thumbnail of the image and the metadata associated with the image. A URL is created for every file in the database so that it can be directly referenced and accessed provided the user has access privileges. Direct file references play an important role within scientific communities, for example the complete URI is used by advanced annotation tools such as Vannotea and the Annotea Side Bar [20]. The assigned Creative Commons License [1] and a link to the complete description, appears below the image.

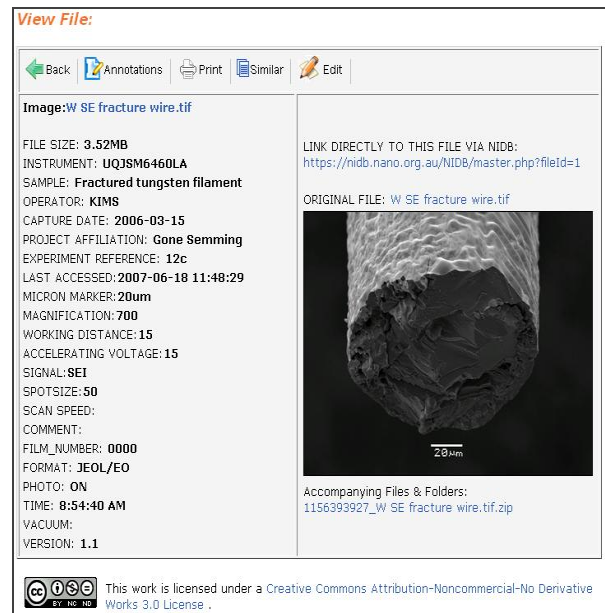


Figure 9: Interface for Viewing Files

Moving Files:

Figure 10 illustrates the interface developed to schedule file movements between the nodes of the grid over AARNET3. The status of all of the user’s file movements can be viewed and refreshed. File movement is required to relocate files to the point of the data processing/analysis.

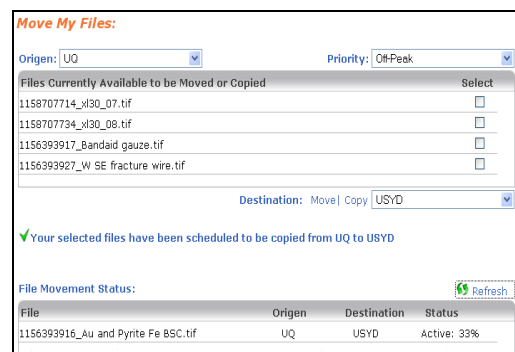


Figure 10: Interface for File Movement

5. Image Processing using Kepler

We have also begun investigating the use of the Kepler workflow system [26] to streamline the image processing and visualization tasks required by NANO users and to more closely integrate the image processing tools with the NIDB and the compute grid through a Web Portal.

Within Kepler, the process of creating a workflow is centered on creating Java classes that extend a built-in Actor class. The existing Kepler release includes a basic *ImageJ* actor which enables processing of a single image using NIH's ImageJ processing library [27]. By using this actor to invoke ImageJ for Microscopy plugins [28], NANO users are able to define workflows that comprise a pipeline of common microscopy image processing operations. However one of the most challenging and compute-intensive tasks facing advanced microscopy today is 3D reconstruction from electron tomography. 3D reconstructions are obtained by processing a series of 2D images captured from a sample tilted at different angles. A typical tomographic data set comprises 151 images taken over an angular range of 150° . In order to speed up the process of 3D reconstruction, ideally the images in the tilt series are segmented in parallel, prior to the image stack alignment step. The existing Kepler *ImageJ* actor can only process a single image at a time. Hence we have developed a new Kepler actor called *ImageJStack* which enables a stack of images to be processed in parallel. We are planning on evaluating its application to 3D cellular tomography at the Institute of Molecular Biology.

6. Telemicroscopy and Annotation

Remote or tele-microscopy enables users at geographically remote locations to access and use specialist instruments without having to travel to the actual instrument. Users can examine their samples under advanced electron microscopes via a real-time high-resolution video streaming interface to the instrument's CCD whilst communicating with a local technician (who is driving the microscope) via video or audio conferencing (e.g., Skype). We have developed a novel web based tool that allows a user to interact with a technician e.g., driving the JEOL6460LA scanning electron microscope at the Centre for Microscopy and Microanalysis, the University of Queensland. The remote session tool, Figure 11, streams high-resolution video footage captured through the secondary electron detector or backscatter detector. We have extended and refined the CyberSTEM system developed previously by the CMM staff [12] and integrated it with the NIDB. Remote users can highlight points of interest to the technician operating the microscope by drawing a "delineator box" on the real-time video display. The

technician can then pan and zoom to the highlighted region and capture images that are instantly updated to the database. This system has been integrated into the portal to provide the scientist with seamless access to the remote instruments, expert technicians, collaboration tools and high resolution images.



Figure 11 : Remote Microscopy with Annotation

7. Conclusions and Future Work

In this paper, we have presented an extensible, scalable, easy-to-deploy framework that will provide the cyber-infrastructure foundations for Australia's expanding characterisation network. In particular GRANI has delivered a distributed image archival and analysis system (with advanced metadata and search capabilities) and collaborative tele-microscopy support services. The long-term archival aspects, developed in partnership with APAC, will prevent loss of data, reduce duplication and facilitate sharing and re-use of research results. The significance of this project is that the system enables more cost-effective, efficient access to and management of the instruments, services and data of the NANO community. The potential impact of the portal is huge, given the range of disciplines, industries and organizations that use NANO facilities. Each of the eight nodes has an average of 300 current clients covering disciplines from nano-materials (novel catalysts and sun-screen materials) to developments in drug delivery, bio-scaffolding and the development of 3D tomographic images of cellular components.

Future plans include: further usability testing and refinement by users from multiple disciplines; expanding the system to new instruments; deploying the image database and archival system across new NANO/AMMRF (Australian Microscopy and Microanalysis Research Facility) nodes; further implementation and evaluation of the Kepler workflow interface; grid-enabling the image processing services to expedite the analysis of large collections of images by distributing computation over the Grid.

References

- [1] GRANI (Grid-enabled Archive of Nanostructural Imagery) Project
<http://www.itee.uq.edu.au/~eresearch/projects/grani/>
- [2] L. Robertson "Overview of the Large Hadron Collider (LHC) Computing Grid", May 2007
http://lcg.web.cern.ch/lcg/documents/LCG_overview_may_07.ppt
- [3] M. Bertero, P. Bonetto, L. Carracciuolo, L. D'Amore, A. Formiconi, M. R. Guarracino, G. Laccetti, A. Murlì, and G. Oliva, "MediGrid: A Medical Imaging Application for Computational Grids," 17th International Symposium on Parallel and Distributed Processing, 2003.
- [4] J. Geddes, S. Lloyd, A. Simpson, M. Rossor, N. Fox, D. Hill, J. V. Hajnal, S. Lawrie, A. McLintosh, and E. Johnstone, "NeuroGrid: using grid technology to advance neuroscience," Computer-Based Medical Systems, 2005.
- [5] M. Brady, D. Gavaghan, A. Simpson, R. Highnam, M. Mulet, "eDiamond: A grid-enabled federated database of annotated mammograms." Grid Computing: Making the Global Infrastructure a Reality. NY: John Wiley & Sons, Ltd., 2003.
- [6] S. Hastings, S. Oster, S. Langella, T. M. Kurc, T. Pan, U. V. Catalyurek, and J. H. Saltz, "A Grid-Based Image Archival and Analysis System," Am Med Inform Assoc, 2005.
- [7] S. Hastings, S. Langella, S. Oster, and J. Saltz, "Distributed Data Management and Integration Framework: The Mobius Projec," presented at Proceedings of the Global Grid Forum 11 (GGF11) Semantic Grid Applications Workshop, 2004.
- [8] Scientific Image Database <http://sidb.sourceforge.net/>
- [9] S. Peltier and M. Ellisman, "The Biomedical Informatics Research Network. The Grid, Blueprint for a New Comput Infrastructure," Philadelphia: Elsevier, vol. ed 2, 2003.
- [10] S. Baru, "Storage Resource Broker (SRB) Reference Manual," Enabling Technologies Group, San Diego Supercomputer Center, La Jolla, CA, August, 1997.
- [11] H. Hochheiser and I. G. Goldberg, "Quasi-Hierarchical, Interactive Navigation of Images and Meta-Data in the Open Microscopy Environment," presented at 2006 IEEE International Symposium on Biomedical Imaging, 2006.
- [12] D. Waddell, J. Drennan, and A. McDowall, "CyberSTEM: Telepresence microscopy from Australia," Royal Microscope Society vol. 35, pp. 157-162, 2000.
- [13] G. Andi, "PHP: supporting the new paradigm of situational and composite web applications". Chicago, IL, USA: ACM Press, 2006.
- [14] ImageMagick, <http://www.imagemagick.org/>
- [15] Globus Toolkit <http://www.globus.org/toolkit/>
- [16] Bannon, D., Chhabra, R., Coddington, P., Cox, D., Crawford, F., Francis, R., Galang, G., Jenkins, G., Rosa, M.L., McMahon, S., Rankine, T., Woodcock, R., Wright, A. (2006): "Experiences with a Grid Gateway Architecture Using Virtual Machines", Proc. of First International Workshop on Virtualization Technology in Distributed Computing (VTDC 2006), Tampa, USA, November.
- [17] M. Lorch, S. Proctor, R. Lepro, D. Kafura, and S. Shah, "First experiences using XACML for access control in distributed systems," Proceedings of the 2003 ACM workshop on XML security, pp. 25-37, 2003.
- [18] D. Spence, N. Geddes, J. Jensen, A. Richards, M. Viljoen, A. Martin, M. Dovey, M. Norman, K. Tang, and A. Trefethen, "ShibGrid: Shibboleth Access for the UK National Grid Service.", eScience 2006, Amsterdam
- [19] Shibboleth Enabled Bridge to Access the National Grid Service (SHEBANGS),
<http://www.mc.manchester.ac.uk/research/shebangs>
- [20] R. Schroeter, J. Hunter, J. Guerin, I. Khan and M. Henderson. "A Synchronous Multimedia Annotation System for Secure Collaboratories" 2nd IEEE International Conference on E-Science and Grid Computing (eScience 2006). Amsterdam, Netherlands. December 2006. p 41. 10.1109/E-SCIENCE.2006.261125
- [21] J. S. Grethe, C. Baru, A. Gupta, M. James, B. Ludaescher, M. E. Martone, P. M. Papadopoulos, S. T. Peltier, A. Rajasekar, and S. Santini, "Biomedical informatics research network: building a national collaboratory to hasten the derivation of new understanding and treatment of disease," Stud Health Technol Inform, vol. 112, pp. 100-9, 2005.
- [22] S. Hastings, T. Kurc, S. Langella, U. Catalyurek, T. Pan, and J. Saltz, "Image processing for the grid: a toolkit for building grid-enabled image processing applications," Proceedings. CCGrid 2003. pp. 36-43, 2003.
- [23] J. Johnston, A. N. A. Hochheiser, and I. G. Goldberg, "A Flexible Framework for Web Interfaces to Image Databases: Supporting User-Defined Ontologies and Links to External Databases," IEEE International Symposium on Biomedical Imaging, 2006
- [24] Creative Commons <http://www.creativecommons.org/>
- [25] J. Bresnahan, M. Link, R. Kettimuthu, D. Fraser and I. Foster. "GridFTP Pipelining" *Proceedings of the 2007 TeraGrid Conference*, June, 2007
- [26] B. Ludäscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger-Frank, M. Jones, E. Lee, J. Tao, Y. Zhao. "Scientific Workflow Management and the Kepler System". *Concurrency and Computation: Practice & Experience*. 2005 Published Online: 13 Dec 2005.
- [27] ImageJ – Image Processing and Analysis in Java
<http://rsb.info.nih.gov/ij/>
- [28] ImageJ for Microscopy – Image Processing and Analysis in Java
<http://www.macbiophotonics.ca/imagej/index.htm>