



University of Michigan: How Grid Science and MGRID are Changing Research and Education



January 2005

NSF Middleware Initiative (NMI) Integration Testbed Case Study Series

Series contact: Mary Fran Yafchak, Southeastern Universities Research Association,
maryfran@sura.org.

The NMI Integration Testbed Program provided practical evaluation of NMI components within the context of real projects and application scenarios from June 2002 through November 2004. During that time, NMI Testbed sites collectively submitted over 220 evaluation reports to middleware component developers as direct feedback into the NMI development cycle. Site representatives also actively inspired, promoted and facilitated the integration of middleware throughout their institutions.

The NMI Integration Testbed Case Study Series documents the most significant influences and outcomes of NMI Testbed sites' middleware integration efforts, highlighting intersections with established projects, application contexts and influences, drivers for innovation, decision points and challenges. Through this documentation, the work of these pioneering institutions is captured to provide a breadth of insight and approaches for others to use towards successful middleware development and deployment.

This NMI Integration Testbed Case Study Series is sponsored by the National Science Foundation Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT) Consortium of EDUCAUSE, Internet2, and SURA. Additional support was provided by the National Science Foundation Cooperative Agreement NSF 02-028, ANI-0123937.

Copyright © 2005 University of Michigan. The University of Michigan permits use of this content for noncommercial purposes with proper attribution. All rights reserved.



Executive Summary

The University of Michigan's MGRID (Michigan Grid Research and Infrastructure Development) project is a pilot institutional grid that is demonstrating the value and utility of grid computing for the University and beyond. Now more than simply a demonstration, MGRID has become a convenient and reliable way to enhance and connect the research activities on campus with those at other campuses across the United States. Dedicated researchers, faculty, and administrators, looking for a grid solution to support their respective research grants with a common campus infrastructure, have collaborated in funding, R&D and staffing to deploy MGRID.

Middleware is a key ingredient in deploying a grid, and MGRID developers have integrated NMI (NSF Middleware Initiative) software into the MGRID environment (e.g., [Condor-G](#)). Where current NMI components have not fully addressed MGRID needs, developers have deployed some tools of their own design (e.g., Walden). As part of their on-going collaborations with grid technologists elsewhere, MGRID developers have made tools such as Walden available to the broader grid computing community. In addition, MGRID developers' collaborations with peer institutions during the NMI Integration Testbed¹ project influenced their

choice of NMI components to be used in MGRID.

The computational problems of many seemingly very different applications on the University's campus can benefit from the same data and computing resources made available to them by MGRID. Determining what these common core needs are and providing for them is the responsibility of MGRID's technical staff, while technical representatives from collaborating application domain projects (e.g., SAKAI, ATLAS) focus on providing for the unique grid requirements specific to their project.

Application domains with projects currently utilizing MGRID resources include particle physics detection, research in the earthquake field, and x-ray crystallography. An important benefit from grid computing common among these diverse application domains is the ability to maximize the "research dollar" by allowing researchers to participate in projects that they would otherwise not be able to work on, or to access remote instruments they might otherwise not be able to access.

resources across the Internet, NMI deployed a practical deployment and evaluation effort called the NMI Integration Testbed. Managed by the Southeastern Universities Research Association (SURA) on behalf of the NMI-EDIT Consortium (NSF Middleware Initiative-Enterprise and Desktop Integration Technologies; see <http://www.nmi-edit.org/>), the Testbed consisted of eight universities that participated in a closely coordinated effort to deploy and evaluate NMI technologies. <http://www1.sura.org/3000/NMI-Testbed.html>

¹ As part of its overall effort to develop and disseminate software that lets scientists and educators share



MGRID resources are offered to users via a prototype portal. The customization of this portal is a prominent example of the University of Michigan's user-centric view of grid technology. The University plans to continue its development of portal technologies. MGRID developers will use these portal developments, as well as grid technologies of their own design and available through national and international

grid forums, to enhance MGRID and bring the benefits of MGRID to additional applications on campus and elsewhere.

For more information about how grid science and MGRID are changing research and education at the University of Michigan, contact Shawn McKee at smckee@umich.edu.



NMI Components Highlighted in this Case Study

The NMI components discussed in this case study series encompass NMI Releases 1 through 4. Information about NMI Releases can be found at <http://nsf-middleware.org/>.

Condor-G

The GRIDS Center's Condor-G is a computation management agent for the grid. Condor-G is the marriage of technologies from the Condor project and the Globus project (see below).

Home site: <http://www.cs.wisc.edu/condor/>; Globus (see below).

Globus

The GRIDS Center's Globus Toolkit is an open-source collection of modular technologies that simplifies collaboration across dynamic, multi-institutional virtual organizations. It includes tools for authentication, scheduling, file transfer and resource description.

Home site: <http://www-unix.globus.org/toolkit/>

KX.509/KCA

NMI-EDIT's KX.509 and KCA provide a bridge between a Kerberos and PKI infrastructure. These tools enable the PKI-based security infrastructure of the Globus Toolkit to integrate with Kerberos-based authentication implemented at university campuses.

Home site: http://www.citi.umich.edu/projects/kerb_pki/

PBS Accounting Toolkit

The GRIDS Center's PBS accounting toolkit takes standard OpenPBS or PBSPro batch scheduler accounting data for Linux and converts it to an XML representation for easier transformation/rendering to other formats. Includes java code, supporting XSD, and XSL files.

Home site: <http://pbsaccounting.sourceforge.net/>

MyProxy

The GRIDS Center's MyProxy is a credential repository for the grid. MyProxy provides a set of flexible authorization mechanisms for controlling access to the repository.

Home site: <http://grid.ncsa.uiuc.edu/myproxy/>

UberFTP

The GRIDS Center's UberFTP is the first interactive, GridFTP-enabled ftp client. It supports GSI authentication, parallel data channels and third party transfers.

Home site: <http://dims.ncsa.uiuc.edu/set/uberftp/>

GridSolve

The GRIDS Center's GridSolve provides the middleware necessary to create a seamless bridge between the simple, standard programming interfaces and desktop Scientific Computing Environments (SCEs). It is a RPC based client/agent/server system that allows one to remotely access both hardware and software components.

Home site: <http://icl.cs.utk.edu/netsolve/>

PERMIS

NMI-EDIT's PERMIS is an authorization infrastructure that uses X.509 attribute certificates (ACs) to hold the credentials assigned to users. PERMIS uses hierarchical Role Based Access Controls, where the X.509 ACs hold a user's roles, and superior roles inherit the privileges of subordinate roles.

Home site: <http://sec.isi.salford.ac.uk/permis/>



University of Michigan: How Grid Science and MGRID are Changing Research and Education

The University of Michigan has been a pioneer in the development of advanced networking technologies, starting with Merit (regional network) in the 1960's. The University continued to break new ground with NSFNET (the academic precursor to the Internet) in the 1980's, and later as a founding member of Internet2¹ in the 1990's. Since early on, researchers and educators at the University have been aware of the need for additional computing power, high-speed networking, and the opportunity to use "idle cycles" of systems not in use to do additional calculations. In recent years, there have been plenty of computers and an ever-increasing number of "clusters" on campus, but no way to determine which were available, how to access them, or how much computing power each one possessed. The many sub-networks carrying data traffic for the University were similarly untapped.

Increasingly today, the ability to harness additional computing power is provided by grid computing. The goal of the Michigan Grid Research and Infrastructure Development (MGRID) project is to develop and deploy a pilot institutional grid to demonstrate the value and utility of grid computing for the University of Michigan and beyond (1). Today, students, faculty and staff at the University of Michigan are able to connect to diverse research resources via a

secure high-speed network. This interactivity has been made possible by the efforts of a small group of dedicated researchers, computer scientists and administrators at the University. Through MGRID, this development team is creating a convenient and reliable way to enhance and connect the research activities on campus with those at other campuses across the United States, enabling access to secure, shared resources while also providing grid resources for other institutional purposes.

This article will show how MGRID has enhanced learning and research opportunities at the University of Michigan by allowing students and faculty to collaborate securely and to perform complex computations on parallel systems located at the University and elsewhere. It will also discuss performance measures and test procedures that University of Michigan staff has developed to support the use of MGRID. Additional focus will be given to the influence of the University's participation in the National Science Foundation (NSF) Middleware Initiative (NMI) Integration Testbed, an evaluation of NMI middleware components within real project contexts including the deployment of MGRID.



The Road to MGRID

Many groups across the University participate in computational-, network- and/or data-intensive research grants. Prior to MGRID, the various departments were working to support their respective research grants using the GRIDS Center's [Globus Toolkit](#)². However, the [Globus Toolkit](#) was difficult to run, and complex to install and manage across these many departments, particularly in the light of the campus's large size. This provided an impetus for the University to seek a grid solution – MGRID – that would allow the departments to develop and support a common infrastructure. MGRID would thus enable more efficient and effective technical support, allowing scientists to spend more time doing research and less time working on IT management problems. It would also improve scalability for a user base of over 175,000 and leverage technical solutions unique to the University of Michigan IT environment, such as the extensive use of Kerberos identities and X.509 entries.

Resources

MGRID was of interest to several units on campus, but none was able to bear the full burden of funding the creation of a national grid node on campus. MGRID therefore became a collaborative unit in funding as well as in R&D. Multiple entities at the University now share the funding of MGRID, including the College of Engineering, the School of Information, the College of Literature, Science and the Arts, and the

Mental Health Research Institute. Additional major funding came from the Office of the Vice President for Research and the Office of the Provost. Initial project funding is for three years, starting in FY'03.

MGRID staffing is also collaborative in nature. After an initial demonstration precursory of MGRID capabilities at iGrid2002 (see "NMI Component Use in MGRID" section), MGRID founders, led by Homer Neal, requested funding from the University to continue working on collaborative technologies. A community of technical staff from various collaborating departments united to create the MGRID architecture and integrate new middleware components effectively. Currently, three full-time MGRID technical staff and several technical representatives from the other participating departments support MGRID core efforts. Much of the main campus networking infrastructure has been updated and improved, catalyzed by MGRID.

Participation in NMI Integration Testbed

Around late 2001, when the University of Michigan started planning and developing MGRID, SURA (Southeastern Universities Research Association) released the Call for Participation in the NMI Integration Testbed project. In this, University of Michigan grid technologists saw an opportunity to perform real life evaluation and provide feedback on NMI middleware software, specifications and services, while at the same time enhancing the work of their faculty and researchers.

Shawn McKee and Victor Wong, the

² Globus information: <http://www-unix.globus.org/toolkit/>



University's NMI Integration Testbed project leads, wanted to incorporate middleware authentication and authorization technologies into their projects and realized there were synergies in initiating the MGRID and NMI Integration Testbed participation simultaneously.

In June 2002, the University of Michigan became the first institution outside of the SURA region to be accepted into the NMI Integration Testbed. Over the two and a half years since then, NMI software releases have been tested and integrated into the MGRID environment with the aim of establishing the requisite middleware for a seamless and scalable grid. It is within this framework that NMI Integration Testbed Program and the MGRID project have informed and benefited each other.

NMI Component Use in MGRID

In September 2002, the first large demonstration precursory of MGRID capabilities took place at iGrid2002, using a real-time video link between Michigan and CERN in Amsterdam for the iGrid conferenceⁱⁱ, an "International Grid applications-driven testbed event...to utilize multi-gigabit experimental optical networks." The University of Michigan collaborators showed a video presentation from Ann Arbor to the conference, loaded the network with other traffic (effectively jamming the video) and was then able to demonstrate implementing QoS (Quality of Service)

prioritization for the interactive video/audio stream over the other traffic. This successful demonstration used a part of [Globus](#) called GARA, as well as the Kerberos-leveraged public key infrastructure. The University of Michigan contributed the latter to the NMI, and it is now known as NMI-EDIT's³ [KX.509/KCA \(Kerberized Certificate Authority⁴\)](#).

Core MGRID infrastructure was developed and demonstrated in April 2004. At present, MGRID core infrastructure connects data and computing resources, making them available for scientific and technical computations at the University to a scheduling mechanism through a portal. The key issues for setting up the MGRID portal were the need for *authentication* and *authorization* and for resource usage *accounting* and *auditing*. These four needs are sometimes known as "AAAA". Grid computing can solve the problem of cross-realm (e.g., inter-institutional) collaboration by using open AAAA standards and established technologies. The portal uses secure authorization based on the University's "Uniqname" identity service and NMI-EDIT's [KX.509/KCA](#) (also known as simply "[KX.509](#)").

The [KX.509](#) security protocol that University of Michigan researcher Bill Doster developed is a key technology for

³ NSF Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT): <http://www.nmi-edit.org/>

⁴ KX.509/KCA information: http://www.citi.umich.edu/projects/kerb_pki/



authentication in MGRID. [KX.509](#) bridges the existing standard for authenticating *within* a community (Kerberos) to the existing standard for authenticating between communities (X.509). By running [KX.509](#) client software, users are able to use their existing Kerberos tickets to authenticate their request to KCA (the [KX.509](#) server) for the signing of a new X.509 client certificate.

NMI Integration Testbed Evaluations

During their participation in the NMI Integration Tested, MGRID developers evaluated a wide range of NMI components, including the GRIDS Center's [Condor-G](#)⁵, [MyProxy](#)⁶, [UberFTP](#)⁷, and [GridSolve](#)⁸.

Condor-G

The University of Michigan looked at both Condor and [Condor-G](#). Condor's strength is in its ability to make "worker node" machines work together as a computing resource. [Condor-G](#), on the other hand, is the marriage of technologies from the Condor and the [Globus](#) projects and makes Condor more "grid-aware". [Condor-G](#) is a computation management agent that provides front-end scheduling functionality for Condor.

The [Condor-G](#) scheduling program proved to be very useful in MGRID applications. Both Condor and [Condor-G](#) are already used extensively for scheduling (the latter as the front-end scheduler) in the ATLAS

project (see ATLAS description in "MGRID Applications" section). However, during their review, MGRID staff found their applications would benefit by having accounting information from Condor processed jobs output in a format compatible with the GRIDS Center's [PBS Accounting Toolkit](#)⁹. MGRID staff will provide a modified Condor scheduling package that creates this type of output for other interested campus projects.

MyProxy

[MyProxy](#) is a credential repository for the grid, allowing users to retrieve proxy credentials as needed, without having to manage multiple certificate files. Initially MGRID developers didn't have use for [MyProxy](#), since it uses short-term, temporary X.509 credentials via [KX.509](#). However, [MyProxy](#) does provide access to other types of long-term X.509 credentials. This feature can be very important for certain projects, like ATLAS, that utilize such credentials. In the future, MGRID would like to provide [KX.509](#) access to [MyProxy](#) servers, allowing transparent access to long-term credentials for [KX.509](#) users.

MGRID developers did note some shortcomings in [MyProxy](#), including the need for a user to manage an additional username/password, and that [MyProxy](#) cannot be used for mutually authenticated SSL since the private key must be available locally.

⁵ Condor-G information: <http://www.cs.wisc.edu/condor/>

⁶ MyProxy information: <http://grid.ncsa.uiuc.edu/myproxy/>

⁷ UberFTP information: <http://dms.ncsa.uiuc.edu/set/uberftp/>

⁸ GridSolve information: <http://icl.cs.utk.edu/netsolve/>

⁹ PBS Accounting Toolkit information: <http://pbsaccounting.sourceforge.net/>



UberFTP

UberFTP is an interactive GridFTP-enabled client. In addition to standard FTP operations such as uploading and downloading documents, **UberFTP** supports Grid Security Infrastructure (GSI) authentication, parallel data channels, striping, and partial retrieves/stores. After some testing, this tool was deemed interesting because of its potential for high-bandwidth transfers of large amounts of data. However, since most I/O is “local” over NFS file systems, MGRID doesn’t require **UberFTP**, at least not yet.

GridSolve

MGRID developers tested **GridSolve**, a client-server system that enables users to solve complex scientific problems remotely. The system allows users to access both hardware and software computational resources distributed across a network. For implementation in MGRID, however, the project team selected MARS, a project under local development that provides a “superset” of **GridSolve** functionality and efficient scheduling of many resources simultaneously (see MARS details in the “MGRID Development” section).

PERMIS

NMI-EDIT’s **PERMIS**¹⁰ was tested by MGRID developers as an authentication choice for their Network Testing and Performance (NTAP) project research on MGRID’s end-to-end network path (see

“NTAP” in “MGRID Development” section). The testing uses a secured framework for authentication and authorization to find failure rates along the network path and to find ways to improve end-to-end performance to overcome sources of non-congestion packet loss. However, MGRID developers chose not to use **PERMIS**, as they determined the University of Michigan developed Walden component would be more useful for the fine-grained authorization required in NTAP. Unlike **PERMIS**, Walden uses an OASIS eXtensible Access Control Markup Language (XACML) policy engine and LDAP (Lightweight Directory Access Protocol) (see Walden in the “MGRID Development” section below).

MGRID Applications and Development

Through MGRID, grid computing is becoming recognized as a solution to many computational problems on campus. MGRID provides the means to support applications and projects that require grid functionality today. Several notable application domain partners are collaborating within MGRID, exchanging requirements and developments to meet increasingly demanding application needs. MGRID is also serving as a proving ground for University of Michigan researchers to design new and extended grid tools and services, influenced by their experience with NMI as well as their existing middleware expertise.

¹⁰ PERMIS information:
<http://sec.isi.salford.ac.uk/permis/>

MGRID's technical staff focus on providing the MGRID core services common among applications, while technical representatives from collaborating application domain projects focus on providing for the unique grid requirements specific to their project. The collaborative staff design is an effective way to tailor the MGRID environment for users while ensuring development resources are used efficiently. Together, this team is developing the MGRID architecture and integrating new middleware components into MGRID.

The overarching goal for is for MGRID to provide for the entire span of needs across all of its application domains. While each application domain is distinct, and on the surface seemingly very different from other MGRID applications, there are in fact a number of common needs among them. These needs include providing common, user-friendly ways for application domain staff to access and use resources on campus and beyond. Security, authentication and authorization issues for users accessing these resources are also high priority needs common among application domains.

MGRID staff is using NMI components and tools of their own design to integrate existing information services on campus with MGRID. Much of their work has been done during the timeframe of the NMI Integration Testbed. The Testbed provided a framework in which MGRID staff could further their analysis of what specific middleware

components would be useful for MGRID. In turn, the analysis MGRID staff have done on NMI components has informed their NMI Integration Testbed collaborators, and the larger grid computing community, about how middleware can be used at other institutions in the US and beyond.

MGRID Applications

Partner applications domain projects within MGRID include ATLAS, NEESGrid, Michigan Center for Biological Information (MCBI) and the Life Sciences Institute. Each of these application domains can benefit from grid computing.

The ATLAS Project – ATLAS (A Toroidal LHC ApparatuS) is a particle physics detector under construction at CERN in Switzerland. Within the US ATLAS Grid, there are many sites with varying numbers of CPUs and various levels of technical ability. Parallel computing jobs reside in the main database where, to achieve the greatest efficiency, “around the clock” human intervention by a staff person is required to submit jobs to each site and manage them. The ATLAS Grid project uses Grid3 tools that are based on the Virtual Data Toolkit that provide services from [Globus](#), Condor and GriPhyN. In the ATLAS project, MonaLisa and Ganglia are used for monitoring, while [Condor-G](#) is used for scheduling.

However, MGRID developers have found that while [Condor-G](#) is intended to assign and manage jobs, current versions aren't



aware enough of conditions (e.g., how many jobs are running; how many processors are available, etc.), there is no real-time feedback, nor enough understanding in the system to recognize site availability, job status and to keep queues filled. To overcome these current limitations in **Condor-G**, MGRID developers may provide and test a metascheduler, such as MARS (see “MGRID Development” below), for use in ATLAS.

NEESGrid – The Network for Earthquake Engineering Simulation (NEES) is an ambitious national program to advance the study of earthquake engineering and to find new ways to reduce the hazard earthquakes represent to life and property. Tom Finholt, the Director of MGRID, helped prepare the NEESGrid User Requirements document that delineates the need for secure access, data storage, high performance computing capacity, open scientific collaboration and remote instrumentation.

MGRID can help maximize the “research dollar” and allow researchers to participate in projects that, without grid computing, they would otherwise not be able to work on. For instance, remote instrumentation allows very expensive tools to be shared among researchers. “Shake tables” used to study earthquakes are a prime example. These tables are huge, room-sized earthquake simulators. The remote instrumentation provided by NEESgrid is giving additional researchers access to these shake tables. NEESgrid staff themselves did a lot work

with middleware components. Their work helped inform the University of Michigan’s evaluation of NMI component use within MGRID by providing a conduit for middleware information flow. NEESgrid staff attended MGRID communication meetings where MGRID’s work with NMI components, and in the NMI Integration Testbed specifically, would be discussed. This collaboration between application domain and MGRID developers working in the NMI Integration Testbed helped MGRID staff customize the MGRID portal (see “MGRID Development” section below) to address the need for secure access to remote instruments for MGRID applications.

Life Sciences: X-Ray Crystallography – Like the shake tables used by earthquake researchers, remote access to other simulators, microscopes, and additional instruments used by researchers working on x-ray crystallography (the x-ray diffraction of molecule structures) can be provided by the use of middleware in MGRID. Researchers in the University of Michigan’s Life Sciences Institute received a grant for access to the beam line at Argonne National Laboratory (ANL) that can use to make x-rays from electrons. To do this, University of Michigan researchers currently send a material sample to ANL, where the Argonne robot puts the samples in a beam line. Sensors read the diffraction pattern of the molecule, and forwards the data back to the university. Integrative middleware will allow the University’s researchers to remotely control this entire process. The existing software at



Argonne and the University of Michigan are not well integrated, so it will be greatly improved by the use of grid computing for remote manipulation and offline processing.

MCBI – The Michigan Center for Biological Information is a nonprofit entity established to build bioinformatics infrastructure and promote bioinformatics research, product development and education in Michigan. MCBI is researching appropriate hardware, middleware, and networking structures for statewide analysis and data-sharing in bioinformatics projects.

MGRID and MCBI have recently started collaborating on a cross-realm bioinformatics portal that will allow researchers at the University of Michigan and Michigan State University to collaborate and share selected computational resources at the two institutions. Secured cross-realm authentication (e.g., using [KX.509](#)) will be one of the important issues for MGRID to address to help the MCBI enhance communications and establish collaborations between academic and industry members.

SAKAI – Concerned with the increasingly high cost of commercial course management systems and possible conflicts with the control of course management and content by proprietary entities, the University of Michigan, Indiana University, Massachusetts Institute of Technology, Stanford University and uPortal undertook leadership of the Sakai project. Over

seventy institutions are currently participating in the Sakai Educational Partners' Program. The project is developing an open source collaboration environment, using portlets to provide academic institutions with a common framework to trade, co-develop, and customize individual applications.

The University of Michigan is making a significant investment to incorporate the Sakai portal application into the campus infrastructure, including integration with MGRID development. The MGRID and Sakai projects share several key staff members, basing this collaboration on the mutual benefit each project will reap through the ability to easily customize portals. Sakai will provide an underlying infrastructure for portals (similar to OGCE) that allows other pieces of software (applets or portlets) to be plugged into it, and it is anticipated that future MGRID portals will incorporate this framework.

These portlets allow the user's portal environment to be customized. Customization of the user portal environment is an integral and important part of the Sakai system. For example, CTools, a Sakai product that combines course, research, and project management into one system, allows faculty to add research results to course content. Several Sakai portlet plug-ins provide functionality unique to the course environment (e.g., calendaring and scheduling functions), however there are a number of generic plug-



in functions that can be adapted for use in MGRID's portal. The emphasis of the prototype MGRID portal (detailed in the next section) is to provide users with secure, user-friendly access to grid computing resources. By adapting Sakai portlets, MGRID developers can further customize the MGRID portal and increase its usability.

MGRID Development

The University of Michigan has leveraged the experience it has in grid computing and grid development for the benefit of both MGRID and the broader grid computing community. The NMI Integration Testbed provided a forum in which MGRID developers could both share experience and learn from other collaborators in the Testbed. Informal project discussions and formal conferences added to MGRID's staff perspectives on which MGRID user needs were unique and which were common within other research communities and institutions.

Where current NMI components have not fully addressed their specific application needs, developers have deployed other tools on MGRID. University of Michigan staff frequently makes these tools, often of their own design or modifications to other products (including NMI components) available to the broader grid computing community. Thus, the MGRID development projects described below may not only benefit MGRID users, but other grid computing users as well.

The Prototype MGRID Portal provides users with lightweight secure access to MGRID-enabled computing clusters on campus through an easy to use GUI interface. The MGRID Portal serves three critical functions:

- Centralizes the installation of heavyweight grid software like [Globus](#) onto the server machine. The user's client machine therefore need only have a browser program (and [KX.509](#)) to make use of the grid software on the server. This means the user and/or technical support staff have much less configuring of the user's client machine to perform.
- Provides security- interactions between the lightweight user client machines and the grid (or Portal) server are secured through the use of the [KX.509](#) component.
- Presents users with a friendly GUI interface of their application's activity choices (e.g., "run a job", "access a remote instrument") and the ability to undertake those activities (run application commands). The portal extracts information from the grid about the applications and resources available to users and presents it to them in an "easy to digest" format.

Without the prototype MGRID Portal, the user must install all of the complex grid software on their machine, plus learn and use arcane "command line" commands to interface with grid resources.



The prototype portal uses CHEF (predecessor to Sakai) to combine OGCE portlets with MGRID portlets, and uses [KX.509](#) authentication as an integrated framework for authorization and authentication. As with Sakai, because the prototype MGRID Portal is based on portlets, it can be customized for different user applications and for specific users (the latter is done by adding a pre-determined search function to the front-end of the portal). Other campus groups will also be able to “brand” their own portals based on the prototype MGRID portal.

In the future, the MGRID portal will have an integrated metascheduler so users won't have to decide where to run their computing job and/or which system is accessible. In addition, through their work in the NMI GridNFS (Grid Network File System) projectⁱⁱⁱ, developers at the University of Michigan Center for Information Technology Integration (CITI) plan to implement efficient file transfers between systems through the deployment of a common MGRID file system.

Network Testing & Performance (NTAP) – CITI at the University of Michigan has been researching the end-to-end network path for MGRID. The principal idea was to attach a tracking system to each router and do point-to-point tests in order to determine which point(s) on the network is/are failing. To do this, they have developed software using [Globus](#) and Kerberos/[KX.509](#) to provide a secured framework for authentication and

authorization. In practice, the network administrator has a list of all users who are authorized to connect to the network. CITI performs their test link by link to find failure rates along the network path. Of the other invited institutions at the NSF network measurement workshop in December 2003, none of the other invited institutions were yet experimenting with this process. Since then, Merit Networks and Internet2 have expressed interest in using the system to test the stability of their own networks.

CITI is also working on improving end-to-end performance to overcome sources of non-congestion packet loss. They are trying to improve throughput fairly, without “slamming” other flows on the network. Currently, when the network's TCP router sees a packet loss, throughput falls to half until the system can recover. The bit error rate – a fundamental coding hardware loss – is high enough for TCP to keep it from ever reaching 10GB on a wide area network (WAN). This means it can take up to 10 hours to recover full packet flow on a 10GB network.

Walden^{iv} – Walden is being used as the MGRID middleware component for providing authentication and authorization services. After determining the [Globus](#) Gridmap file wasn't able to scale to a large university setting, MGRID staff developed and implemented Walden, which integrates [Globus](#) with existing campus mechanisms. To accomplish this, Walden's developers designed it to take advantage of a new



callout feature in [Globus 2.4](#) that allows replacement of [Globus's](#) static Gridmap file. Walden replaces this file with an XACML policy engine and LDAP, allowing MGRID's access policies to be implemented^v (2) Walden allows for guest accounts, multiple resource ownership, virtual organizations and can handle multiple gatekeepers. More information about Walden development and availability may be obtained from Beth Kirschner, MGRID, at bkirschn@umich.edu.

MARS (Michigan Advanced Resource Scheduler) – Within a distributed environment, grid users and administrators often need to determine the optimal and differentiated scheduling of resources such as clusters, network bandwidth (e.g., for video-conferencing) and data resources. MARS, a grid-level meta-scheduler of clusters or other resources, can be used to determine which machines (or nodes) are available given a set of jobs. MARS was developed to fill a need to provide this higher-level scheduling that was not present in the existing NMI components. The metascheduler allows differentiated priorities for submitted jobs (e.g., on-demand and best-effort). It can provide on-demand jobs precedence over other jobs running on computational cluster(s) via job preemption and advanced reservation. With the beta version of MARS currently operational, the next step is to integrate it with Walden (described above) so that fine-grained authorization and differentiated scheduling can be provided simultaneously. A paper^{vi} by MGRID staff about MARS was accepted

for the 5th ACM/IEEE International Workshop on Grid Computing conference (held in November 2004, in conjunction with SC04). In the future, MARS could also be used for scheduling MCUs (Multipoint Control Units) for wide-area multi-institution videoconferencing. More information about MARS may be obtained from Abhijit Bose, MGRID, at abose@umich.edu.

The Road Ahead

The strategic scope for MGRID has changed from a focus on traditional high performance computing (HPC) to a broader vision focused on large collaboration environments, computing, middleware, grid technologies and cyber-infrastructure. NEESgrid is a good indication of what can be expected from grid-enabled portals – infrastructure should become invisible to end users so that these researchers can spend more of their time on scientific tasks rather than managing cyber-infrastructure. In the next year, the goal of MGRID is to facilitate this user-centric view of grid technology and cyber-infrastructure. MGRID will not only continue to participate in NMI, it will also partner with Sakai to create opportunities for broader acceptance of Sakai (including a grid-enabled version) within the collaborative research and instructional environment at the University of Michigan. Further, MGRID will continue to work on grid- and cyber-infrastructure projects, such as GridNFS, Walden, and MARS, and to participate in national and international grid application projects, such as ATLAS and NEESGrid.



More Information

For more information about how grid science and MGRID are changing research and education at the University of Michigan, contact Shawn McKee at smckee@umich.edu.

References

- (1) <http://www.mgrid.umich.edu>
- (2) Kirschner, et al. *Walden: A Scalable Solution for Grid Account Management*, <http://www-personal.engin.umich.edu/~hacker/papers/gridmapfile.pdf>

Information in this paper is also based on interviews with the following individuals, on August 3, 2004:

- Dr. Victor Wong, MGRID executive committee member, Director of CITI, ykw@umich.edu
- Dr. Shawn McKee, Assistant Research Scientist, Physics, MGRID, smckee@umich.edu
- Dr. Tom Finholt, Director of MGRID, finholt@umich.edu
- Dr. Abhijit Bose, Associate Director of MGRID (incoming) abose@umich.edu
- Dr. Tom Hacker, Assistant Director of MGRID (outgoing) hacker@umich.edu
- Beth Kirschner, Systems Research Programmer III, MGRID bkirschn@umich.edu
- Eric Hofer, Computer Systems Consultant III, School of Information, ehofer@umich.edu



Links of Interest

GRIDS Center <http://www.grids-center.org/>

iGrid2002 www.igrid2002.org

MGRID www.mgrid.umich.edu

MARS <http://www-personal.engin.umich.edu/~abose/website/mars.htm>

NEESgrid www.neesgrid.org/

Network Testing & Performance (NTAP) www.citi.umich.edu/projects/ntap/

NMI-EDIT <http://www.nmi-edit.org/>

NMI Integration Testbed Program <http://www1.sura.org/3000/NMI-Testbed.html>

NSF Middleware Initiative <http://www.nsf-middleware.org/>

Sakai sakaiproject.org/

University of Michigan www.umich.edu

University of Michigan Center for Information Technology Integration (CITI) www.citi.umich.edu/

University of Michigan Life Sciences Institute

<http://www.lifesciences.umich.edu/institute/index.html>

ⁱThe Internet2 Project is a consortium being led by 206 universities working in partnership with industry and government to develop and deploy advanced network applications and technologies, accelerating the creation of tomorrow's Internet. <http://www.internet2.edu/about/>

ⁱⁱ www.igrid2.org

ⁱⁱⁱ GridNFS: Grid Network File System <http://www.citi.umich.edu/projects/gridnfs/GridNFS-SoW.pdf>

^{iv} The word "Walden" is a derivation of warden, meaning "keeper of".

^v Walden moves user authentication to the client and replaces the static mapping between X.509 identities (Distinguished Names) and local user names in the *Globus* grid-mapfile with a dynamic approach based on *kx509* and secure LDAP. More information about Walden can be found in the paper: Walden: A Scalable Solution for Grid Account Management, Kirschner, et al, <http://www-personal.engin.umich.edu/~hacker/papers/gridmapfile.pdf>

^{vi} More information on MARS can be found in the paper "MARS: A Metascheduler for Distributed Resources in Campus Grids" by A. Bose, B. Wickman and C. Wood, 5th IEEE/ACM International Symposium on Grid Computing (Grid2004), October 8, 2004, Pittsburgh: <http://www-personal.engin.umich.edu/~abose/website/marsgrid2004.pdf>