

NMI Integration Testbed Program – Final Report

NSF ANI-0123937, “Supporting Research and Collaboration through Integrated Middleware”

M.F.Yafchak, 6/15/05

Materials noted as archived on the SURA Web site are available at: <http://www1.sura.org/3000/NMI-Testbed.html>

The NMI Integration Testbed program, funded through the NMI-EDIT consortium, brought real life feedback to bear on the design and evolution of NMI middleware from September 2001 through May 2005. A middleware-enabled, robust IT infrastructure is key to growing the science capacity of the South and the nation - a primary focus of the SURA mission. Through the NMI, SURA members and collaborators contributed to the immediate advancement of scientific applications and the establishment of a sustainable cyberinfrastructure for broad and future use.

Program Startup, Structure & Reporting

NMI Testbed sites were selected in March 2002 through an open Call for Participation and external review process (<http://www1.sura.org/3000/testbedcfp.pdf>). The CFP and Testbed site selection was conducted under the review of the NMI Management Team and in such a way as to ensure that the Testbed program addressed the evaluation of both EDIT and GRIDS components. SURA led the development of several versions of the CFP in order to meet this goal.

NMI Testbed site participants' qualifications, diversity and collaborative spirit were key to the program's success. Each institution was selected for their readiness for active middleware integration as well as how their expertise would complement the balance of experience and perspective in the Testbed overall. Each site leveraged institutional resources to complement NMI resources and advance middleware integration in both research and enterprise applications, uncovering commonalities in how such integration could succeed.

Sponsored Sites (under contract to SURA and receiving NSF funds to complete their work): University of Alabama at Birmingham, University of Florida, Texas Advanced Computing Consortium (University of Texas at Austin), University of Virginia.

Un-sponsored Sites (under contract to SURA but contributing their own resources to participate): University of Alabama at Huntsville, Florida State University, Georgia State University, University of Michigan.

The University of Southern California also participated informally but was not under contract.

From June 2002 through August 2002, communication and reporting processes were initiated within the Testbed in parallel to verification of NMI Release 1. Milestones included establishment of intra-Testbed communication mechanisms, development of a Testbed Council as liaison to the NMI Management team, establishment of a SURA NMI Advisory Group comprised of SURA member CIOs, and debut of the public NMI Testbed Web site.

Program progress was reported through contributions to NMI-EDIT annual reports to the NSF, SURA annual reports to UCAID, and periodic status reports made available on the NMI Testbed Web site. Annual and status reports are archived on the SURA Web site.

All NMI Testbed sites under contract to SURA submitted annual performance reports to SURA as part of contracted deliverables. This final report combines information from site reports with observations and documentation of the NMI Testbed Manager and serves as the overall program report. NMI Testbed Site reports are on file at SURA, 1201 New York Ave, NW, Suite 430, Washington, DC, 20005.

NMI Component Evaluation

NMI Integration Testbed sites submitted over 220 evaluation reports to NMI component developers as feedback into the NMI development cycle. In the second year of the NMI – the first year of NMI Testbed operation – NMI Testbed sites completed evaluation of the first and second NMI releases. Over 60 evaluation reports were completed for NMI Release 1, covering all 18 components with a minimum of 2 evaluation reports per component. Fifty-nine reports were completed for NMI Release 2, covering each component with at least one evaluation and an average of three per component. In the following year, Testbed sites completed evaluation of NMI Releases 3 and 4, collectively producing 104 additional reports. In the final 6 months of NMI Testbed site participation, SURA and NMI Management agreed that sites' limited time and resources would be better spent on documentation of middleware integration vs. the evaluation of NMI Release 5.

NMI Testbed sites committed to the evaluation of various components in each release based on their areas of expertise, overlap with institutional objectives, potential for integration with real projects and applications, and, on occasion, recommendations or requests from the NMI Management team. Feedback was provided according to Component Testing Guidelines that were developed through collaboration of the NMI Testbed Manager and the Testbed Council, a sub-set of the NMI EDIT and GRIDS Center teams. In this, the Testbed Council communicated on behalf of the Management team and relevant developers to circumscribe evaluation priorities, depth, and direction. The Guidelines were tailored to each release and provided per-component detail on evaluation objectives, support and final report deadlines. Completed evaluation reports were sent in their entirety to the NMI Testbed Manager and to the Testbed Council who forwarded reports to applicable development teams.

As the NMI moved from Release 1 to Release 4, there was increased emphasis from the NMI Management team and the NSF to identify concrete examples of value being achieved through the integration of NMI middleware. The evaluation tasks and reporting formats of the Component Testing Guidelines evolved accordingly and the trend among Testbed sites' evaluation overall was towards more practical critique born of deeper project integration. Copies of all component evaluation reports submitted by the NMI Testbed sites are included on a CD provided with this report. Component Testing Guidelines are archived on the SURA Web site.

NMI Middleware Integration

Throughout the NMI Integration Testbed program, NMI Testbed sites leveraged component evaluation to bolster access to existing resources and offer new applications and services:

Nearly all NMI Testbed sites have initiated, catalyzed or solidified their deployment of key research and enterprise infrastructure through their participation in the NMI. Seven sites have extended their enterprise authentication, authorization and directory services to provide centralized information for applications such as click-to-dial video conferencing, self-service email configuration, campus white pages,

authenticated downloads for IT management, and integrated course management systems. Seven also have active efforts underway to build campus-wide research grids that will be based on or are heavily integrating NMI components.

NMI Testbed sites leveraged their experience evaluating NMI components to extend accessibility within other research projects. Examples include:

- Facilitated access for University of Michigan researchers' to the ATLAS Grid Testbed (<http://gate.hep.anl.gov/gfg/usatlasgrid/>)
- University of Alabama in Huntsvilles' development of grid-based collaborative applications for the NCSA MEAD expedition (<http://www.ncsa.uiuc.edu/expeditions/MEAD>), NASA Marshall Space Flight Center's Space-based Science Operations Grid, and the LEAD (Linked Environments for Atmospheric Discovery) project.
- University of Virginia's extension of the large-scale molecular dynamics code, CHARMM, to operate on an NPACI Grid using Globus, greatly increasing the number of users who might use CHARMM in a Grid environment.
- Georgia State University's communication with physicists on campus towards facilitating their ability to participate in the PHENIX/RHIC (Relativistic Heavy Ion Collider) project (Dr. Xiaochun He, <http://petitt.phy-astr.gsu.edu/welcome.html>) and the Particle Physics Data Grid (<http://www.ppdg.net/>).

Georgia State University initiated a project to bring nuclear physics experiences to the desktops of local area Georgia high schools. This project is being led by Dr. Xiaochun He and related to the "Quarknet" project (<http://quarknet.fnal.gov/>), of which GSU is a 2003 Center. The project is incorporating GRIDS technology and leveraging Georgia State's NMI Testbed experience to manage distributed muon detectors in high schools throughout the state, enabling students to participate directly in physics research.

Several NMI Testbed sites were involved in independent middleware development projects and integrated their experiences and perspectives from the Testbed into this work. Among these are:

- University of Alabama at Birmingham (ViDe.Net Middleware Project, <http://lab.ac.uab.edu/vnet/>)
- Georgia State University ("Semantic Facilitator™ (SM)" - display and analysis of directory objects; "Muon Detector Grid: Wide Area Implementation")
- University of Michigan ("Middleware for Grid Portal Development, <http://www.collab-ogce.org/nmi/index.jsp>)
- Texas Advanced Computing Center (the NPACI Grid Portals project, <https://hotpage.paci.org>); "Middleware for Grid Portal Development", <http://www.collab-ogce.org/nmi/index.jsp>)
- University of Virginia ("Complex Application Requirements in Grid Computing/Legion-G/OGSA.NET).

As noted in NMI Testbed site plans, the participation of PIs for these projects within the NMI Testbed is providing a valuable means of reciprocal feedback.

Results Dissemination and Broadening the Community

Perhaps the most significant contribution of the program was outreach to new or extended user communities, disseminating experience from a variety of implementations to meet campus needs:

The program team produced four full-day workshops to disseminate evaluation and integration

results. These events were held April 2003 in conjunction with the Internet 2 Spring Members' Meeting, November 2003 preceding the EDUCAUSE 2003 Annual Conference, September 2004 following the Internet 2 Fall Members' Meeting and as a full-day preconference session at the EDUCAUSE 2004 Annual Conference in October 2004. All NMI Testbed sites participated in these events and spoke in-depth on their experience and lessons learned while integrating NMI middleware into campus grids, research grids, and enterprise-wide identity management and related applications. Workshops were relatively small but well received and interactive, drawing nearly 150 attendees overall – a mix of technical implementers, IT management and research faculty, including some international attendees. Agenda and presentations from these workshops are archived on the SURA Web site.

The NMI Testbed Manager and NMI Testbed sites also provided materials and presentations at numerous other academic events, including EDUCAUSE Annual Conference 2002, 2003, 2004; Internet Spring and Fall Members' meetings 2002, 2003, 2004; GlobusWorld 2004, and via the AccessGrid; SuperComputing 2004; Oklahoma SuperComputing Symposium 2004. Additionally, Testbed site representatives proactively pursued dissemination opportunities on their own through presentations at other conference venues and in internal meetings and training events at their institutions.

In Spring/Summer 2003, SURA led the development and delivery of a training workshop on the deployment of NMI Directory Services, in collaboration with four Testbed sites (GSU, UFL, UAB, UAH) and NMI-EDIT Outreach. Participants' experience was leveraged to prototype training for the implementation of NMI-based Directory Services for under-represented communities. With funding from the NCSA Alliance PACS (Partnership for Advanced Computational Services), workshops were delivered to early adopters from the Consortium of Independent Colleges and the NSF AN-MSI (Advanced Networking, Minority-Serving Institutions). The content and instructional methods emphasized a process-oriented approach to enabling implementation and demonstrated the value of such practical workshops to extend middleware deployment. Materials from this workshop are available on the SURA Web site at <http://www1.sura.org/3000/NMIIndex.html>.

NMI Testbed sites and collaborators convened in January 2004 to share institutional experiences in grid technology and identify synergistic activities that could be valuable on a regional or extra-regional scale. Grid technology is an important element in the SURA IT Strategy and also in plans to advance academic and research infrastructure at the many academic research institutions, including NMI Testbed sites. The group explored how the expertise developing in the NMI Testbed could be applied to the benefit of future collaborative work and identified project directions that would benefit grid development at large. Preliminary "thinking groups" formed to refine select ideas and recommend next steps. Results to date include initiation of the SURA Cyberinfrastructure Workshop Series, new associations for collaborative proposal development, and formalization of a SURA HPC-Grid Initiatives Planning Group to support regional grid development.

REU (Research Experiences for Undergraduates) Supplement

SURA further expanded NMI awareness through proposal and receipt of an NSF REU (Research Experiences for Undergraduates) Supplement to the NMI Testbed award in Spring 2003. Five REU positions were awarded at four institutions - Georgia State University, Texas Advanced Computing Center, University of Michigan and the University of Virginia. These positions were grounded in the qualifications and experience of the host institutions and unified through the processes of the NMI

Integration Testbed Program. Students evaluated and deployed NMI components in a variety of settings, helping to inform development while increasing their understanding of middleware and its potential to change the way scientific research is done.

Georgia State University students worked on two specific projects, “Muon particle detector GRID for K-12” and “GRID-enabled Applications/Graphic Rendering for the GSU Digital Aquarium”, with the overall objective of creating a process for grid-enabling applications at GSU. Both students were also awarded scholarships to attend a one-week intensive Summer Workshop in Grid computing and its application in scientific data analysis (http://cgwa.phys.utb.edu/outreach/gw_data_workshop/index.php). The University of Michigan experience focused on exploring various NMI components to create a new integrated middleware package for the ATLAS project. TACC and UVA students collaborated with NMI team members in grid portal research and deployment in conjunction with the NMI Testbed Grid.

REU progress reports were submitted from REU sites to SURA at the half way and final points of each REU experience. Culminating REU reporting from REU sites was included in their NMI Testbed Program final reports. Progress reports and related materials are on file in the SURA office, 1201 New York Avenue NW, Suite 430, Washington, DC, 20005. Related excerpts from NMI Testbed Site final reports are included as Attachment B, with budget data summarized with the overall project budget in Attachment A.

NMI Testbed Grid Project

Throughout the duration of the NMI Integration Testbed program, interest and activity towards the deployment of grid technology increased at nearly all sites. In some cases, this was in response to Testbed site representatives increasing their own familiarity with grid tools, recognizing the value this technology can bring to campus researchers, and seeking to raise the awareness and use of grids at their institution. In other cases, Testbed site representatives themselves are expert in grid research and development and identified critical areas for grid evolution that the Testbed was uniquely positioned to address. As experiences and perspectives were shared, sites expressed a desire to work together within the virtual organization of the NMI Testbed to support and grow this activity.

In light of this, SURA formalized a program for intra-testbed grid deployment and requested approval from UCAID for a redirection of unused hardware funds from the SURA NMI budget to provide some support for the effort. The project formally began in September 2003, with the following objectives:

- Provide a mechanism for Testbed sites to explore grid capability with researchers and faculty at their institution and with others considering grids as a collaborative technology.
- Enable the identification, documentation and resolution of unaddressed cross-campus authentication issues necessary to support dynamic, multi-project (vs. project-specific) inter-institutional grids.
- Expand evaluation and real life usage of NMI components related to “plumbing” grid authentication to the campus enterprise infrastructure.
- Position the Testbed to make additional valuable contributions to the NMI with existing funding and provide project participants with a foundation from which to plan and fund additional work.

Georgia State University, an unsponsored Testbed Site, was sub-contracted to lead immediate deployment and the cataloging of applications using or having the potential to benefit from grids. This work was

supported by the reallocation of existing NMI Testbed funding: 17.3K initially reallocated from unused hardware funds in the general program budget, and an additional 28.2K from funding returned to the project by the University of Florida (see Attachment A). The University of Virginia, a sponsored Testbed site, led the examination of issues to be resolved for scalable inter-institutional authentication leveraging campus enterprise credentials, and the development of the NMI Testbed Bridge Certificate Authority (Bridge CA). This effort was supported through a redirection of resources currently being applied to component evaluation at UVA. Products from this include the Catalog of Grid Applications (http://art12.gsu.edu:8080/grid_cat/index5.jsp), deployment of the Bridge CA (<https://www.pki.virginia.edu/nmi-bridge/>), Bridge CA concepts as an area of grid research (<http://www.cs.virginia.edu/~humphrey/papers/BridgeCAGridSecWorkshop2004.pdf>) and establishment of an ongoing commitment to create an inter-institutional, multi-purpose grid. This work is documented in the NMI Integration Testbed case study “Exploring Technical and Policy Considerations for Inter-Institutional Grids” (see next section).

The NMI Testbed Grid effort has since evolved to become SURAgrid, a multi-institutional effort to model scalable heterogeneous (vs. project-specific) inter-institutional resource sharing. SURAgrid will evolve based on standard, open protocols to enable distributed resources to be shared securely across institutional boundaries without sacrificing local autonomy. As of this writing, eighteen institutions are participating, including the original NMI Testbed participants, additional SURA members and several collaborators from outside the SURA membership. Participants have committed an initial set of resources to be shared via SURAgrid and meet regularly, facilitated by SURA, to evolve a persistent grid and grow the application set. SURA is also providing funds through March 2006 for work in key areas (application development, Bridge CA implementation, portal access, communications) while the project seeks external sources of support. For more information on SURAgrid, see <http://www1.sura.org/3000/SURAGrid.html>.

NMI Integration Testbed Case Study Series

As sites proceeded with their middleware integration, it was clear that valuable experience was being gained through the process of middleware integration in addition to the data requested through component evaluation reports. SURA discussed this with the Testbed Council mid-way through the project to find ways to capture this material. There were no resources available to support this documentation in either the NMI-EDIT (including the Testbed) or GRIDS center programs, however, there was strong agreement among all (including the NSF) that such documentation would be of notable value to NMI outreach. Several months later, when funding became available from the return of a portion of site funding from the University of Florida (see Attachment A), SURA requested and was approved to reallocate 68K of this funding to create the NMI Integration Testbed Case Study Series. NMI Management and the NSF agreed that the final months of Testbed Sites’ energy and resources would be better spent assisting in further articulating and documenting these case studies rather than on evaluation of NMI Release 5.

The NMI Integration Testbed Case Study Series captures some the most significant outcomes of the program’s integration efforts and provides a breadth of insight and approaches for others looking to deploy. Cases documented represent middleware integration at various stages of implementation and also the influence and impact of NMI component evaluation on the development of other middleware products and services.

NMI Integration Testbed Case Study Series

Collaborative within the Testbed program

- Exploring Technical and Policy Considerations for Inter-Institutional Grids: NMI Testbed Grid
 - Technical Supplement: Authentication & Authorization in SURAGrid: Concepts and Technologies
 - Technical Supplement: Building a Campus Grid: Concepts and Technologies (*in progress - estimated completion June 30*)
- Research Experiences for Undergraduates within the NMI Integration Testbed Program

Focus on enterprise infrastructure

- University of Alabama at Birmingham: Blazing a Path - Directory Services Evolution at UAB (*pending completion of final review by UAB and Testbed Council*)
- University of Florida: Identifiers, Social Security Numbers, and Identity Management
- Florida State University: Identity Management & NMI-Component Integration
- Georgia State University: Building an Identity Management Infrastructure for the eUniversity
- University of Southern California: Shibboleth and Pubcookie at USC- Authentication and Authorization for All
- University of Virginia: Campus PKI Services/Bridge CA

Focus on grid infrastructure

- University of Alabama in Huntsville: Data Mining in the Modeling Environment for Atmospheric Discovery
- University of Alabama in Huntsville: Grid Integration with LEAD (Linked Environment for Atmospheric Discovery)
- Georgia State University: NMI Grid Middleware Enhances Georgia State's Cosmic Ray Muon Project
- Georgia State University: Shifting the Grid Paradigm with NMI
- University of Michigan: How Grid Science and MGRID are Changing Research and Education
- Texas Advanced Computing Center at The University of Texas at Austin: UT Grid Design & Implementation
- University of Virginia: Protein Folding on the Grid

Some Lessons Learned

Funding Testbed Site activity at this level did not prove to be key to performance. Instead, critical success factors were the alignment with institutional objectives (amount of success was perhaps most affected by the amount of alignment that could be achieved), self-motivation and collaborative nature of the site representative(s), alignment with the site representative(s) normal work expectations and responsibilities, and the reality of time. A close working relationship between IT and research was also very important to making progress on the integration of middleware into existing projects and applications, and to advancing the understanding of middleware as an underlying foundation for both enterprise and research IT.

In cases where funding was provided to Testbed Sites, the amount was limited and therefore expected to defray rather than cover costs. Given that, SURA's approach of providing general guidelines for the use of funds but leave final spending decisions to the sites proved to be very effective. Sites could choose to supplement institutional resources as needed to meet deliverables. SURA's guideline that approximately 2/3 of the funding be spent to insure staff participation also proved valuable, with most sites electing to spend the bulk of funding on staff, particularly after the first year.

The process of actually integrating middleware into an existing service or application takes much longer than expected, with areas affecting integration being political and cultural as much as technical. Lots of time (often no substitute...) and communication are required to build an understanding of middleware and its benefits before integration can proceed. In the NMI Testbed, this resulted in less "real life" evaluation than intended but much was learned about the path/process for middleware integration. In addition, "lab testing" proved more valuable than expected as it was used to build necessary experience on the part of site representatives in order to match component potential to institutional objectives and gain buy-in

from the application/project “owners”. Some formal connections or forums between the leads of desirable projects and the larger NMI program could assist with this and also help insure that developing middleware is immediately useful even as it evolves towards being more broadly applicable. The complex nature of the middleware integration process also affected the output of the NMI Integration Testbed Case Study Series, with the early estimate of 25 to 40 case studies at 4 to 6 pages each being realized as 16 case studies at 10 to 25 pages each.

Many people will work on NMI “for NMI's sake” - to help bring about the transparent and secure information and resource sharing that is the vision for middleware. There was real desire among Testbed Sites to coordinate and collaborate in order to catalyze this reality. In addition, sites found that the association with a national effort had a very positive impact on their ability to engage others in middleware integration. This validated and supported them in their role as “middleware champions”.

Don't underestimate the value of involving “fresh” perspectives. Constructing the NMI Testbed from a Call for Participation rather than assembling a group of pre-existing collaborators brought new viewpoints and expanded involvement to the NMI. Requiring both research and IT to be represented on Testbed Site teams brought melded experience to bear on both evaluation and integration. The involvement of students – undergraduate students through the REU supplement and graduate students involved at several sites – enabled more progress to be made with input from diverse perspectives and experiences.

On the specific structure and processes of the Testbed:

- The rapid release schedule adopted by the NMI is more conducive to lab testing and staging than integration and use within actual projects. It was very difficult for Testbed Site representatives to keep up with changes on all components in use, verify components each time they changed, and convince application or project leads to upgrade every six months.
- “Feedback on the feedback” is very important. Site evaluation reports were routed directly to NMI developers but there was no commitment on the part of the developers to provide feedback in return on if or how this input was used. Sometimes this was obvious from changes (or not) in the next release and, on occasion communications with developers did take place, particularly once facilitated by the Testbed Council after this issue was illuminated. However, there was no regular or guaranteed feedback and, in the absence of that, some sites struggled to remain motivated and produce consistently comprehensive reports.
- People are the most important resources for doing this type of work and people working well together is key to success. The role of the Testbed Manager was critical to establishing a collaborative atmosphere and insuring the completion of Testbed Site deliverables. SURA would recommend that the NMI program continue to support and emphasize the involvement of staff focused on coordination, motivation, and mechanisms to bring, and keep, teams working together.

Attachment A: Preliminary (Un-audited) Actual versus Budget

The information below is based on preliminary information from SURA Accounting for the period September 1, 2001 through the middle of May 2005, and detail from the Testbed Manager's records related to budget reallocations throughout the project. A final budget report will be sent from SURA Accounting to UCAID, Attn: Lisa Hogeboom, following final close out of figures through May 31, 2005.

The initial program budget included the following categories and amounts per year:

Senior Personnel - \$25,000 salary + \$6,250 fringe

Equipment, (Hardware & Software) - \$5,825

Domestic Travel - \$8,000

Sub-awards - \$240,000

Meeting Support - \$10,000

Indirect Costs @ 10% - \$4,950

TOTAL = \$300,025

Between August and November 2003, the University of Florida fell behind in meeting its contracted obligations due to personnel issues. UFL initiated a change in PIs and worked to rectify the situation but, after several months, this had only limited success. SURA then placed a financial hold on NMI funding to UFL, requesting receipt of specific overdue deliverables or termination of the NMI Testbed site agreement. UFL chose to continue and worked with SURA to develop a modified Statement of Work but also returned 96.6K in funding that they determined could not be used within the time frame of the project. UFL eventually completed most of the adjusted Statement of Work but did not submit additional invoices, a final Testbed Site report to SURA, or provide an explanation to be included with this report, despite repeated communication and escalation from SURA,.

In April 2004, SURA requested and was approved to reallocate the returned funds from UFL to supplement the Testbed Grid project (28.2K sub-award to GSU) and to develop the NMI Integration Testbed case study series (45K writer salaries, 9K travel, 14K publication & dissemination). This also included a no-cost project extension through February 2005.

In February 2005, SURA requested and received an additional no-cost extension through May 31, 2005 to complete the case study project, which included a reallocation of remaining funds:

Funds remaining as of February 28, as estimated at time of request: 33.5K

Reallocation:

Senior Personnel - 20K

Travel – 11.8K

Other – 1.7K

Total = 33.5K

Note that this did not include or address the category of indirect funds. It was determined by SURA Accounting that SURA could not invoice for indirect but the decision was not formally confirmed in time for the Testbed Manager to consider a reallocation to program activities.

Preliminary Actual vs. Budget – SURA Detail

Categories do not reflect all changes to budget allocations over time. The final report from SURA Accounting will provide that information.

| | NMI expenses through 5/31/05 | NMI Budget | Difference |
|--------------------|-------------------------------------|----------------------|---------------------|
| Senior Personnel | \$ 88,258.10 | \$ 75,000.00 | \$ (13,258.10) |
| Fringe Benefits | \$ 22,064.53 | \$ 18,750.00 | \$ (3,314.53) |
| Equipment | \$ - | \$ - | \$ - |
| Domestic Travel | \$ 36,019.41 | \$ 33,000.00 | \$ (3,019.41) |
| Case Study Writers | \$ 48,882.50 | \$ 45,000.00 | \$ (3,882.50) |
| Other | \$ 35,938.97 | \$ 53,930.28 | \$ 17,991.31 |
| Indirect | \$ - | \$ 22,675.00 | \$ 22,675.00 |
| Total | \$ 231,163.51 | \$ 248,355.28 | \$ 17,191.77 |

Preliminary Actual vs. Budget – Sub-Award Detail

| Institution | NMI expenses through 5/31/05 | NMI Budget | Difference |
|----------------------------|-------------------------------------|----------------------|---------------------|
| Univ of Florida | \$ 27,567.11 | \$ 74,030.17 | \$ 46,463.06 |
| Univ of Virginia | \$ 179,868.49 | \$ 180,000.00 | \$ 131.51 |
| Univ of Alabama Birmingham | \$ 180,000.00 | \$ 180,000.00 | \$ 0 |
| TACC/Univ Texas at Austin | \$ 176,181.05 | \$ 180,000.00 | \$ 3,818.95 |
| REU & Testbed Grid - GSU | \$ 58,707.81 | \$ 59,154.55 | \$ 446.74 |
| REU - Univ of Mich | \$ 4,692.75 | \$ 6,820.00 | \$ 2,127.25 |
| REU – TACC/UT Austin | \$ 6,304.51 | \$ 6,820.00 | \$ 515.49 |
| REU – Univ of Virginia | \$ 6,225.88 | \$ 6,820.00 | \$ 594.12 |
| Total | \$ 639,547.60 | \$ 693,644.72 | \$ 54,097.12 |

Notes:

UFL – See overall project history above. NMI Budget column reflects amount remaining after return of 96.6K.

TACC – From TACC re unspent funds: “The unspent funding for the SURA Testbed site consisted of mostly fringe benefits and a small amount for international travel. Fringe benefits are estimated at 25% and the actual expended amount was less based on specific staff members appointed to the effort. In addition, no foreign travel was necessary to complete the project during this time period. The unspent REU funds were also associated with costs associated with WCI, UCI and OASI. The student hired required less fringe than the budgeted estimate. In addition, the student hired worked a partial month in June.”

GSU – Accounting reports currently bundle GSU Testbed Grid and REU expenses. Total NMI expenses above reflect figures from the latest Accounting report plus a final incoming invoice from GSU, as added by the Testbed Manager.

UMich – From UMich re: unspent funds: “Constraints on the students from classes and exams reduced the amount of time they had to devote to the REU program to approximately 2/3rds of the expected time commitment. In addition our first student, Karen, was unable to continue working on the project after classes ended at the end of April 2004 and we were unable to locate a replacement student until Fall 2004 when Richard started. Also budgeted was money to support a trip to a relevant conference or meeting (\$650), which went unused because Karen left the project earlier than anticipated and Richard had insufficient time to consolidate his work for presentation.”

Attachment B: REU Supplement Final Report Data

The following are excerpts are from NMI Testbed Site final reports, submitted January 31, 2005)

Georgia State University

Georgia State received funding (\$12,000 + 1,640 Travel) for NSF REU students with funding received September 2003. We engaged two students beginning August 2003, Nicole Geiger, Physics, and Anish Shindore, Computer Information Systems. Funding extension (Amendment # 4) as well as of supplements from Georgia State's Advanced Campus Services have enabled them to remain engaged in the project throughout.

Their work developed the Catalog of Grid Applications to include over 475 researchers from 20 universities and national labs. (See: http://art12.gsu.edu:8080/grid_cat/index5.jsp)

Both attended the SPI GRID Summer Workshop, June 21-25, 2004 as scholarship students and earned "high marks" from their being so well prepared as a result of their NMI Integration Testbed experience. Both attended the Internet2 Fall Member Meeting, September 27-October 1, 2004, Austin, Texas and were contributing participants in Georgia State University's demonstration of OGCE Portal, Ganglia monitoring, My-Proxy certificates, globus-run to set user environment and submit a job, grid-ftp to transfer results to a laptop grid node, where results were visualized in 3D model.

Participated in deployment of grid-enabled PC clusters for Cosmic Ray Lab and the ACS development/test grid.

University of Michigan

UM was funded for one Research Experience for Undergraduate (REU) award with Shawn McKee as the REU mentor. During the first part of the REU award we hired Karen Hayrapetyan, who started October 2003 and finished in May 2004. His experience was documented in the "NMI Integration Testbed: Research Experiences for Undergraduates within the NMI Integration Testbed" report as well as Karen's write-up of his experience. Karen summarized his overall experience in the following way:

"I have enjoyed my time that I have spent working on NMI. I have learned a lot about ATLAS project and how information is processed in the scientific world. My experience has included how information is divided through network, how network weather service works, learning about operating systems like Linux and Unix, etc. How the security works, how to protect information. In my opinion the idea of issuing and using certificates is highly convenient and secure method for scientists. Overall the whole NMI project is very important and necessary project. Many things still have to be done especially Graphical User Interface must be developed. I can't wait until everything will be finished and it will be possible to see how final programs work and use them. My project is over but I was introduced to a new world for me and I will continue to explore and keep me up to date even on my own."

While Karen wasn't able to complete an independent project integrating NMI with ATLAS software because of a heavy class load, we feel he benefited greatly from his exposure to middleware and the ever-increasing importance it plays in the conduct of research and education. It was also enlightening for Shawn and those working in MGRID to see a fresh perspective on middleware. The questions and interaction, which happened because of Karen's participation in the REU program, were perhaps as valuable for those of use mentoring Karen as for Karen himself.

Because Karen wasn't able to devote the amount of effort we had planned for to the REU program we had unspent funds in May 2004 and asked for an extension to hire another REU student. Since most students are not around during the summer, we weren't able to locate a suitable participant until September 2004 when we hired Richard French, an aerospace undergraduate with a strong interest in computers and software. Richard completed his REU work in January 2005 and [submitted his final report]. Richard's comments on his experience are contained in his report and we are including a relevant paragraph here which summarizing his experience:

"My experience with the NMI software suite has been very positive. I have started from a position of nearly no Linux and grid computing knowledge to one that is continually growing. At the beginning of my work, which included learning basic Linux operation and the abstract ideas behind grid computing and the ATLAS computing model, I was overwhelmed with the amount of information. By the time I began deploying Release 5.1 I felt a disconnect between what I had read before and what I was reading now; until I found Dr. Wilkinson's documentation. With a fresh look at even the basic prerequisites I began reworking and found I could be much more successful if I kept the reasons for my work in mind: I was configuring the software to accomplish some task. With a renewed sense of purpose I have continued to deploy components of the NMI Release 5.1 with much greater success. Challenges, though, still abound. Dr. Wilkinson's documentation is tailored to his class and his grid. A more general and more detailed combination of the kind he has begun would be extremely beneficial to new grid installers everywhere. A few examples in the NMI Release 5 Installation Guide, written by Jeffery House (one of Dr. Wilkinson's students), will illustrate the power of the method."

Overall the REU program has been very successful from Michigan's perspective. Having undergraduates involved in learning about middleware and, in our case, the ATLAS experiment has been an interesting and enriching experience. We were able to make some significant contributions in debugging NMI releases and providing feedback for NMI documentation. We know that both Karen and Richard enjoyed their participation and gained a significant amount of experience and understanding about both middleware and the role it plays in a scientific collaboration like ATLAS.

Texas Advanced Computing Center

Over the course of the NMI Testbed project, TACC participated in several sub-projects with other member institutions. In September 2003, SURA established the NMI Testbed Grid as a sub-project of the NMI Integration Testbed Program to allow the NMI Testbed members to collaborate on grid research and deployment activities and to seek funding for additional research projects. TACC also hired an undergraduate student under the auspices of the NSF REU program to work on Grid User portals in the context of this Testbed Grid. In collaboration with SURA technical writers, TACC created a case study document on UT Grid, a campus grid project at the University of Texas led by TACC which uses several NMI components. Each of these sub-projects is described in greater detail in this section...

NSF REU Program

During the period from June 1st to August 24th 2004, TACC participated in the National Science Foundation (NSF) Research Experiences for Undergraduates (REU) program by hiring Ashesh Sahib, an undergraduate student at the University of Texas at Austin, to work on grid portal projects related to the NMI Integration Testbed Grid. This REU experience is part of a multi-site project involving three

NMI Integration Testbed sites: TACC, the University of Virginia and the University of Michigan. It was based on work by researchers from each institution in the area of Grid Portals, Open Grid Services Architectures, and Grid Technologies and Security. Students located at each site learned to work collaboratively with each other in order to develop the portal and infrastructure. Additionally, the REU students had the opportunity to work with state-of-the art technologies in web portals and grid technologies (e.g., Jetspeed/Portlet, web services, Globus, GridFTP).

Sahib's main responsibility was to work with other portal developers and with application users on the NMI Integration Testbed Grid to consult on and assist with the development of a User Portal. This portal would provide users with resource monitoring and job management capabilities. Sahib initially spent time becoming familiar with grid concepts and technologies by installing and using the NMI GRIDS components. This was a learning experience, since Sahib had no prior experience in distributed computing. Sahib then used the Open Grid Computing Environments (OGCE [1]) software to create a user portal with capabilities for a grid user to monitor grid resources and to submit and manage the execution of jobs on these grid resources. As part of the effort of building the user portal, Sahib had to update the resource monitoring scripts that provided information to the User Portal. During this time, Sahib participated in meetings with the TACC portal team, and with Art Vandenberg's NMI Integration Testbed team at Georgia State University.

One of the features lacking in the OGCE toolkit is the ability for a grid user to specify a series of tasks and have the grid middleware manage the workflow on their behalf. Although several workflow tools exist, none of them has been integrated with user portal software such as OGCE. As a research component in his work, Sahib investigated Pegasus [2], a grid workflow management tool, and the feasibility of integrating it into a user portal toolkit. Pegasus was developed at the Information Sciences Institute at the University of Southern California, and uses several NMI components, such as the Globus Toolkit and Condor-G.

University of Virginia

As part of the NMI SURA program, TACC took the lead and established a research agenda for undergraduates focused on Portal security. One student was at TACC and one student was at the University of Virginia. At the University of Virginia, Jonathan Martin was a third-year student who worked with PI Humphrey on Portal Security during the summer of 2004. He specifically focused on Shibboleth, MyProxy, and PubCookie.

Jonathan's work culminated in a document:

" Extending Existing Campus Trust Relationships to the Grid through the Integration of Pubcookie and MyProxy", by Jonathan Martin, Jim Basney, and Marty Humphrey. Submitted for Publication.

<http://www.cs.virginia.edu/~humphrey/papers/PubcookieMyproxy.pdf>