8B.6

LEAD AT THE UNIDATA WORKSHOP: DEMONSTRATING DEMOCRATIZATION OF NWP CAPABILITIES

Thomas Baltzer, Anne Wilson, Mohan Ramamurthy
Unidata Program Center, UCAR, Boulder, CO

Suresh Marru, Marcus Christie, Dennis Gannon
Indiana University

Al Rossi, Shawn Hampton, Jay Alameda
National Center for Supercomputing Applications (NCSA)
University of Illinois at Urbana-Champaign

Kelvin Droegemeier
Oklahoma University

1. INTRODUCTION

During the week of July 10, 2006 Unidata held its triannual User's Workshop. The theme of this year's workshop was "Expanding the Use of Models as Educational Tools in the Atmospheric & Related Sciences" (See http://www.unidata.ucar.edu/community/2006 workshop/ for more information.) Workshop organizers and LEAD team members saw this as an opportunity to unveil LEAD to and receive feedback from a self-selected group of Unidata community members that are particularly interested in being able to run forecast models. Thus on July 13th 2006,the LEAD lab day, members of the Unidata community got their first experience with capabilities being developed under the Linked Environments for Atmospheric Discovery (LEAD) project. The key LEAD goal demonstrated during the workshop was that of "Democratization," that is, providing capabilities to the larger meteorological community.

The specific capability the LEAD team brought was the ability to run the WRF model over a meso-scale domain selected by each workshop participant.

*Corresponding author address: Thomas S.V. Baltzer PO Box 3000 Boulder Colorado, 80307: tbaltzer@ucar.edu

Democratization was demonstrated in that it typically takes a grad student or new researcher approximately 6 months to begin making runs with a model like WRF. Workshop participants were accomplishing this in a matter of minutes.

At workshop, participants worked with software that demonstrated the specific concepts of:

1) Lowering the barrier to entry by making it easy for users to:
    a. Experiment using meteorological tools
    b. Create meteorological forecasts
    c. Perform mesoscale modeling and forecasting
    d. Access data (both data used by and generated by LEAD)
    e. Make use of large scale cyberinfrastructure (TeraGrid)

2) Giving users freedom from technological issues such as:
    a. Hassle-free access to supercomputing resources
    b. Hassle-free execution of forecast models and related tools
    c. Data format independence

The stated goals of the LEAD Lab day at the workshop were to introduce LEAD concepts

and software to the users and to receive feedback from them, and to describe the education initiatives spawned by LEAD. In addition, the LEAD development team had the goal of performing its successful end-to-end large-scale test in a real world user setting.

2. DESCRIPTION OF THE EFFORT

For the purposes of the LEAD Lab at the workshop, LEAD developers targeted a goal of allowing two sets of 25 users to simultaneously launch a high resolution, steered WRF forecast in which users could select the forecast start time, grid spacing, and spatial domain location. The three options provided were a 5km and 20km horizontal grid spacing for a 600km by 600km domain and a 20km horizontal grid spacing for a CONUS (277km X 176km) domain. For model initialization users could select between the 40km NAM or a 10km spatial resolution ADAS assimilation, while boundary conditions were limited to 40km NAM forecasts.

There were 50 workshop machines available for use. Because each user would be launching two jobs (one from each orchestration system), LEAD Lab organizers planned two separate sessions of 25 each, in order to distribute the computational load. Thus the schedule for the day included two separate morning sessions, each with half the workshop participants, during which users configured their accounts, configured their forecasts, and launched them. Intermingled with these morning sessions was a session on activities being conducted by the LEAD education thrust group. In the afternoon, users came together to view their results.

2.1 WRF Orchestration Systems

The team presented two orchestration systems. One is a Grid Process Execution Language (GPEL) based workflow execution system with a web based portal interface primarily developed at Indiana University (IU), and the second is a desktop client-based system called “Siege” which is based on the Ensemble Broker capabilities developed at NCSA.

Both systems made use of a tremendous amount of infrastructure provided by LEAD and other institutions. The infrastructure common to both orchestration systems used included:

- TeraGrid resources allocated at NCSA
- TeraGrid software stack (primarily GridFTP and GRAM)
- The Weather and Research Forecast (WRF) model software
- Data delivered via the Unidata Internet Data Distribution (IDD) capability
- Storage of IDD data on the Unidata LEAD test bed
- Cataloging and accessing of that data via the Unidata THREDDS Data Server (TDS)
- Visualization of the initial and boundary condition input datasets and the output WRF data was accomplished using the Unidata Integrated Data Viewer (IDV).

2.1.1 Portal and GPEL Orchestration

The Portal and GPEL based orchestration additionally made use of:

- The LEAD test bed hardware system housed at Indiana University
- Front end portal frameworks with a host of portlets using GridSphere and OGCE (Open Grid Computing Environments) technologies
- Metadata crosswalk software that translated THREDDS catalog metadata into the LEAD Metadata Schema
- The LEAD resource catalog and myLEAD catalog and their associated Query Service
- The myLEAD agent interacting with the myLEAD server which is built on top of OGSA-DAI
- The Grid Process Execution Language (GPEL) workflow engine which is a BPEL (Business Process Execution Language) compliant workflow engine
- The Generic Factory Service
- Transient Application Services
- Geographic Region Search Tool (Geo-GUI)
- WS-Messenger, (Notification Broker)
- Host Selection Broker
- Xbaya Workflow Composer

- Dynamic Service Creator
- The ARPS Data Assimilation System (ADAS) for initial and boundary conditions for the WRF run

A snapshot of the user experience in the GPEL solution is shown here:

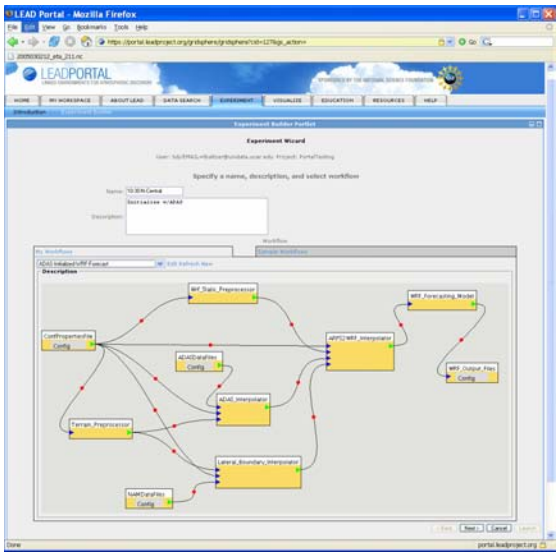

2.1.2 Siege Based Orchestration

The Siege based orchestration additionally made use of:

- The Siege desktop client customized from its native ensemble run orchestration to the single steered WRF run capability
- NCSA's "troll" family of services (ensemble broker, execution service, host information service, notification service, and elf/ogrescript local execution engine)
- A customized IDV plugin that allowed selection of the model domain over the top of the NAM data being used for initial and boundary conditions
- WRFSI for setting up initial and boundary conditions for the model run

A snapshot of the User Experience with the Siege orchestration is shown here:

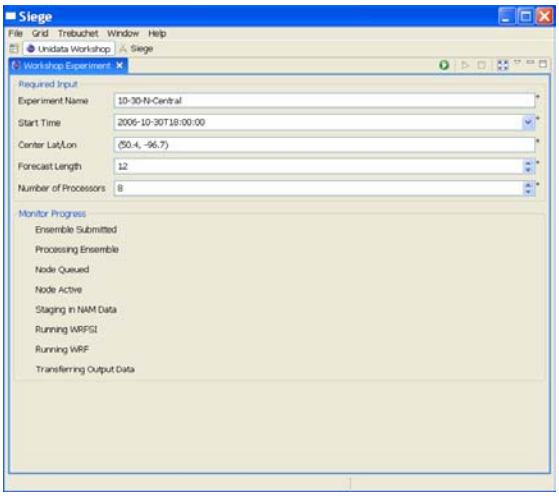

2.2 Workshop Preparation

A significant amount of software development and integration needed to occur prior to the workshop. The LEAD team had developed many of the components involved in the workshop experience, but many had not been integrated and needed additional work to facilitate that integration.

Given the magnitude of integration this represented, in preparation for the workshop, LEAD team members organized three scheduled days of testing where NCSA TeraGrid resources were reserved specifically for LEAD testing on May 24, June 7, and July 6. We were particularly interested in using these days as an opportunity to perform scalability tests on the two orchestration systems. However, these plans were significantly impacted by side effects due to TeraGrid software upgrades that occurred on the resources allocated at NCSA.

These changes and other upgrades and system administration tasks conducted on the NCSA TeraGrid machines had impact upon our preparation and demonstrated a need for the LEAD team to determine how to work in better concert with TeraGrid.

During the workshop, participants were asked to fill out a survey of information detailing their experiences and how they might be able to use the capabilities going forward. This information is seen as a valuable resource to the LEAD team going forward in development and maintenance of LEAD.

In addition, the Michigan Community Engagement team conducted a thorough usability evaluation of the portal environment and services provided through the portal. In addition to redesigning the "look and feel" of the portal, the team suggested human-computer interaction and layout changes to the web services which were largely implemented. To provide consistent and easily navigable information about the project, the entire contents of the LEAD website were redesigned and enhanced to fit within the portal interface.

Given the magnitude of the integration effort and the issues encountered along the path of integration, the entire LEAD team was delighted when we arrived at the workshop with both orchestration systems ready for use. We did not have the chance to perform full scalability testing which did wind up being an Achilles heel for the team.

3. WORKSHOP EXPERIENCE

The LEAD Lab was executed as planned, with two morning sessions to launch forecasts and one combined afternoon session to view results. Also, at the start of the LEAD Lab day, the goals and results of the Community Engagement survey were presented to the workshop participants. During an interactive discussion of several of the survey questions, participants raised the same key issues that were highlighted in the survey, *prior* to learning about the analyzed results. Specifically, participants identified the same time-consuming research tasks (data collection and management), data acquisition priorities (data accuracy and availability of specific variables), and software adoption issues (ease-of-use, new functionality, and long-term support). These observations reinforce the appropriateness of LEAD's future goals, as influenced by the survey results.

Here is an image of workshop participants having a hands-on experience with LEAD.

3.1 Portal/GPEL Orchestration

The portal allowed users to select from a list of four pre-composed workflows. Users first had the opportunity to select a workflow using either ADAS or NAM data for initialization using the ADAS tool for initialization. For each of these workflows they could pick either CONUS Domain, or a regional 600kmX600km domain with either the horizontal grid spacing of 5km or 20km and users were given an option to perform either a six- or twelve-hour forecast over the selected region.

A total of 73 GPEL workflows were launched by the two groups of participants. Four of the workflows failed to finish because the system did not use a comprehensive user input validation, and the four users had selected a wrong combination of ADAS Initialized data and NAM Boundary Conditions.

There were other non-user generated problems. The gridftp server on the Unidata LEAD test bed machine where the input data was hosted and the output results were stored, encountered problems when 10 gridftp concurrent transfers were attempted to/from the machine. The Generic Factory Service crashed twice and one of the Resource Catalog crawls returned no results, the reasons are still unknown at the time of writing this report and are still being investigated.

These problems gave a very good opportunity to resurrect user launched workflows by the GPEL workflow system administrator. The resurrection capability of the workflow system helped to identify the potential failures at run-time and in all cases the second attempt workflow execution finished successfully.

Except the four workflows launched by users with incorrect inputs, the workflows produced meaningful results that were cataloged in each user's personal space accessible from the portal interface, served via the TDS and viewable with IDV.

3.2 Siege Orchestration

Leveraging the general purpose framework for creating workflows Siege provides, workshop participants were presented with a tailored application for launching NAM initialized WRF forecasts using the WRFSI initialization system. In order to give users the complete experience, users downloaded and installed the client software. Next, users could configure a default forecast length and also the number of processors to use. These values can also be overridden on a per-experiment basis.

The first workshop session achieved eighteen submissions of which fifteen were left pending while three became active. This first attempt at a large-scale test showed that Siege services were paralyzed by the volume of events, which provoked a failure in the message bus (ActiveMQ). When messages stopped flowing, subsequent operations were hung as well; this meant that nothing could go forward.

Siege developers determined this problem on the fly and applied a quick patch to the execution service that would mitigate this particular bottleneck by reducing the number of file events. Unfortunately, this required a reboot of all the services. The loss of events meant that even the three active jobs which completed would not be visible to the users who submitted them.

The second session had seventeen submissions (eleven successes and six failures). One failure was due to a user submitting a 24-hour forecast so the job exceeded the time allotment. The rest of the failures were due to missing files, likely due to gridftp problems described next.

A significant overall point of failure for both systems was the the gridftp server on the Unidata test bed, which simply could not serve such a large and concentrated volume of data requests. This in part owing to the 100Base-T network card over which the data was being transferred, however this warrants further analysis. We can improve gridFTP performance via a redesign of storage structure and replacing the 100Mb/sec network card with a Gigabit Ethernet card. But we also need to work with gridFTP developers to provide greater tolerance of a lot of connections, all performing transactions, setting up transfers, etc – but a consistent (and perhaps consistently degraded) level of performance, based on the number of connections against the potential bandwidth. What we saw consisted of unusual and misleading failure conditions and messages that we believe indicated that the gridftp server could not handle the request, but told us something like "not in gridmapfile" (when in fact, they were). We've seen the same behavior to mass store, but worse – in which no error is thrown, but the connections hang (and timeout on the server side, apparently) after about 3-4 concurrent connections are established.

3.3 Viewing Results

In the afternoon session, users had the opportunity to view the results of their WRF runs as well as the data resources that went into their runs and even compare the results with collocated radar observations.

In the case of Portal/GPEL orchestration, users returned to their accounts on the portal and found the URL referencing their WRF product. Clicking on that URL brought the result up in the IDV. Users could then search for the NAM data used for initial and/or boundary conditions as well as ADAS data they may have used for initial conditions and visualize all these in an integrated way. To compare with radar data, users needed to access the data via the TDS interface to these data as they are not yet available in the LEAD Resource Catalog.

In the case of the Siege orchestration, users brought the Siege client back up on their desktop and found the result of their run via that interface. Clicking on the result link brings up in the IDV. Data used for initial and boundary conditions was either already available (as the users had used IDV to select their mesoscale region) or could be obtained

via the TDS. Likewise, users could bring up radar data and other observations to compare with their model results.

3.4 User Experience

The amount of material to be presented was greater than the time allotment provided. Users would have been better served if they could have had more time to assimilate the information and processes presented. To a large extent this problem was unavoidable as the schedule had to be finalized months before it was clear what exactly we would be able to demonstrate at the workshop.

The overall impression from the users was that the capabilities provided by the LEAD team were interesting and useful even though the technologies were still experimental and buggy. Users seemed to appreciate the value of being able to launch a WRF forecast without having to deal with the vast array of technical issues involved.

Perhaps not surprisingly, the user interface is a highly important component of both orchestration systems. In the evaluations, nearly all users had comments about the interface. Some found it easy, some did not, and others had opinions in between. All indicated that more and better documentation was needed. The Michigan Community Engagement team was present at the workshop to observe the use of the orchestration systems and identify further usability improvements. The team identified ways that the functionality and navigation could be clarified and improved in the future, including documentation, based on the participants' difficulties and questions. The team also recorded explicit suggestions for ways of making the software consistent with existing conventions and expectations of meteorologists and students (e.g., preferred date and time notation systems). Some users preferred the web-based approach and indicated that they thought that people and students in particular would be more comfortable with that. Others cited network difficulties as a benefit of the local client approach taken by Siege.

Some quotes from workshop participants:

"Sharing data/models, democratization of data and models were the highlight of the workshop."

"I liked that it was a step by step process. The steps were well broken down into easily understood components."

"Easy to use."

I liked the "ease of running WRF and fact that it might use someone else's computer to do the computing"

"I will implement it in my introductory level Met. and Climate class and workshops for 6-12 teachers"

Overall, the goals of the workshop were met. Users were introduced to the idea remotely launching forecasts and responded in a generally positive manner. Also, developers got plenty of valuable feedback regarding the user experience and stress test results.

4.0 Future Directions

Now that the LEAD team has created a first cut of capabilities for the community we are broadening our exposure through a Beta test program facilitated through the Unidata Program Center. (Please send e-mail to support-lead@unidata.ucar.edu to express interest).

From the feedback provided by workshop and Beta test participants, we have established the next set of requirements we intend to fulfill through development (better data searching and interaction, end-to-end ADAS workflows, improved interfacing in general, etc.) A new release with these capabilities is scheduled for December 2006.

We are also keen on following through with the idea of having LEAD taught in introductory Meteorology classes as alluded to by one of the workshop participants quotes above.

5.0 ACKNOWLEDGEMENTS

LEAD is a Large Information Technology Research (ITR) Grant funded by the National Science Foundation under the following

Cooperative Agreements: ATM-0331594 (University of Oklahoma), ATM-0331591 (Colorado State University), ATM-0331574 (Millersville University), ATM-0331480 (Indiana University), ATM-0331579 (University of Alabama in Huntsville), ATM03-31586 (Howard University), ATM-0331587 (University Corporation for Atmospheric Research), and ATM-0331578 (University of Illinois at Urbana-Champaign, with a sub-contract to the University of North Carolina).

The LEAD project involves the work of nearly 100 individuals whose dedication has resulted in the capabilities that will be shown here. The authors would like to recognize all of them, but in particular we'd like to recognize: John Caron, Rich Clark, Ethan Davis, Charles Hart, Yuan Ho, Scott Jenson, Rob Kambic, Brian Kelly, Ning Liu, Jeff McWhirter, Don Murray, Beth Plale, Rahul Ramachandran, Yogesh Simmhan, Kevin Thomas, Nithya Vijayakumar, Yunheng Wang, Dan Weber, and Bob Wilhelmson.

6. REFERENCES

Droegemeier, K.K. and Co-Authors, 2005: Service-oriented environments in research and education for dynamically interacting with mesoscale weather. *Computing in Science and Engineering*, **7**, 12-29.

Droegemeier, K. K. and Co-Authors, 2005: Linked environments for atmospheric discovery (LEAD): Architecture, technology roadmap and deployment strategy. Preprints, *21st Conf. on Interactive Info. Processing Systems for Meteorology, Oceanography, and Hydrology*, San Diego, CA, Amer. Meteor. Soc.

Droegemeier, K.K. and Co-Authors, 2005: Service-oriented environments in research and education for dynamically interacting with mesoscale weather. *Computing in Science and Engineering*, **7**, 12-29.

Gannon, D. and co-authors: 2005, Building Grid Portal Applications from a Web-Service Component Architecture, *Proceedings of the IEEE. Special Issue on Grid Technology.*

Gannon, D., B. Plale, M. Christie, L. Fang,Y. Huang, S. Jensen, G. Kandaswamy, S. Marru, S. Lee Pallickara, S. Shirasuna, Y. Simmhan, A. Slominski, and Y. Sun, Service Oriented Architectures for Science Gateways on Grid Systems, International Conference on Service Oriented Computing 2005, B. Benatallah, F. Casati, P. Traverso(Eds.), LNCS 3826, pp. 21-32, 2005. Springer-Verlag Berlin Heidelberg 2005. (Invited paper.)

Lawrence, K.A. and I-h Kim, 2006: Survey of mesoscale meteorology community and requirements traceability matrix for LEAD project. Unpublished technical report version 1.0, 38pp. [Available from the Collaboratory for Research on Electronic Work, University of Michigan School of Information, 1075 Beal Avenue, Ann Arbor, MI, 48109-2112 or from the authors at kathla@umich.edu or hwanie@umich.edu.]

Michalakes, J., S. Chen, J. Dudhia, L. Hart, J. Klemp, J. Middlecoff, and W. Skamarock, 2000: Development of a next-

generation regional weather research and forecast model. Proceedings, *Ninth ECMWF Workshop on the use of Parallel Processors in Meteorology*, Reading, U.K., November 13-16, 2000. Argonne National Laboratory preprint ANL/MCS-P868-0101.

Murray, D., J. McWhirter, S. Wier, and S. Emmerson, 2003: The Integrated Data Viewer—a web-enabled application for scientific analysis and visualization. Preprints, *19th Conf. On Integrated Information and Processing*, 8-13 February, Long Beach, CA, Amer. Meteor. Soc.

Plale, B., J. Alameda, R. Wilhelmson, D. Gannon, S. Hampton, A. Rossi, and K.K. Droegemeier, 2004: User-oriented active management of scientific data with my LEAD. *IEEE Internet Computing*, **9**, 27-34.

THREDDS Catalog Specification http://www.unidata.ucar.edu/projects/THREDDS/tech/catalog/InvCatalogSpec.html

OPeNDAP Project http://www.opendap.org/

The Integrated Data Viewer (IDV) http://www.unidata.ucar.edu/software/idv/

The Weather Research & Forecasting (WRF) Model http://www.wrf-model.org/index.php