6B.2

A NEW APPROACH FOR USING WEB SERVICES, GRIDS AND VIRTUAL ORGANIZATIONS IN MESOSCALE METEOROLOGICAL RESEARCH

[*1,2]Kelvin K. Droegemeier, [3]Beth Plale, [4]Mohan Ramamurthy, and [5]Craig Mattocks
[1]Center for Analysis and Prediction of Storms and [2]School of Meteorology
University of Oklahoma
Norman, Oklahoma
[3]Indiana University
Bloomington, Indiana
[4]University Corporation for Atmospheric Research
Boulder, Colorado
[5]University of North Carolina
Chapel Hill, North Carolina

1. INTRODUCTION

Linked Environments for Atmospheric Discovery (LEAD; Plale et al. 2004, 2006; Droegemeier et al. 2005, 2007; Droegemeier 2008) is a National Science Foundation (NSF) Large Information Technology Research (ITR) grant that has pioneered new approaches for integrating complex weather data, assimilation, modeling, mining, and cyberinfrastructure systems. LEAD empowers researchers and students with capabilities heretofore available at only a few major universities and research or operational centers around the world, and it does so using a service-oriented architecture similar in many respects to the familiar Amazon.com, where storage and compute resources are provided via "the cloud." By managing the complexity of inter-operative cyber tools and providing flexibility and ease in how they can be linked, LEAD allows users to focus their time on solving the science and engineering problems at hand, providing a means for more deeply understanding the tools and techniques being applied rather than the nuances of data formats, communication protocols, and job execution environments.

Foundational to LEAD is the notion that today's static environments for observing, predicting, and educating about mesoscale weather are fundamentally inconsistent with the manner in which such weather actually occurs, namely, with often unpredictably rapid onset and evolution, heterogeneity, and spatial and temporal intermittency. To mitigate this inconsistency, LEAD has created an integrated, scalable framework in which meteorological analysis tools, forecast models, and data repositories can operate as dynamically adaptive, on-demand, grid-enabled systems, thus allowing them to a) change configuration rapidly and automatically in response to weather; b) respond to decision-driven inputs from users; c) initiate other processes automatically; and d) steer remote observing technologies to optimize data collection for the problem at hand.

Having completed five years of research (now in a one-year no-cost extension period) including the development and testing of a prototype system as well as operational testing in collaboration with the NOAA National Severe Storms Laboratory and Storm Prediction Center (e.g., Weiss et al. 2007; Kong et al. 2008; Xue et al. 2008), LEAD has advanced research in cyberinfrastructure and virtual organizations through the following community-enhancing contributions:

- As a ***virtual organization*** of atmospheric researchers and computer scientists, LEAD enables productivity gains by: eliminating the need for an individual to write high performance computing (HPC) allocation proposals, accomplished through a single community account to national computing resources on the Teragrid; providing a researcher single sign-on to distributed HPC resources through GSI security; and simplifying access to HPC resources through the LEAD science gateway portal.

- As a ***framework***, LEAD supports analysis pipelines (i.e., workflows) that can contain a mix of community and custom models

[*] Corresponding Author Address: Prof. Kelvin K. Droegemeier, University of Oklahoma, 120 David L. Boren Boulevard, Suite 5900, Norman, OK 73072. Email: kkd@ou.edu. Visit LEAD on the web at http://portal.leadproject.org.

and analysis and assimilation routines, executing on-demand, dynamic forecasts. The framework captures provenance about data objects, the latter being critical to improved reproducibility and sharing of results.

LEAD is examining strategies for deployment as a formal community facility. Such a facility, funded by multiple organizations and made openly available, would support research, education, and operational experimentation including, in particular, the emerging concept of "warn on forecast" (e.g., Burgess et al. 2005; Stumpf et al. 2007; Kuhlman et al. 2008).

We describe herein these plans as well as new efforts to link LEAD technologies with those of the Weather Research and Forecast (WRF) model Portal (e.g., Govett and Smith, 2008).

2. OVERVIEW OF RECENT ENHANCEMENTS

Foremost among LEAD goals during the past year was the ability for users to edit namelist input files and incorporate them into workflows for key applications such as the WRF model and ARPS Data Assimilation System (ADAS). Two approaches were taken, the first involving an extension of existing parameter specification capabilities within the LEAD Portal itself (Figure 2.1), and the second involving a separate but linked web interface for namelist editing and associated file management (Figure 2.2). Both have provision for consistency checking at various levels of sophistication, ranging from simple tests to ensure linear stability in the choice of the WRF timestep to much more complex tests that interrogate choices of interrelated parameters such as mixing coefficients, inputs to physics schemes, etc. Owing to its limited resources, LEAD has chosen to develop a basic checking framework which will allow others in the community to contribute as needs arise. Further details about namelist editing and linkages with the WRF Portal are provided in §4.

In an effort to broaden support for community components, the user can, through the LEAD service factory, convert applications into services for use in LEAD workflows. Through a user interface users can now compose workflows consisting of their own model or analysis components or can edit or use an existing community workflow. Options for repeating workflows at specified times or time intervals, and for pre-scheduling workflows, are available and have been added to the Portal, as well as the important capability of terminating or restarting an executing workflow. Currently, workflows execute on selected TeraGrid resources under a community account. Given sufficient community interest, workflow submission could be enhanced to allow a user to specify a particular computational resource, or ask the system to broker the necessary resources based upon specific quality of service requirements.

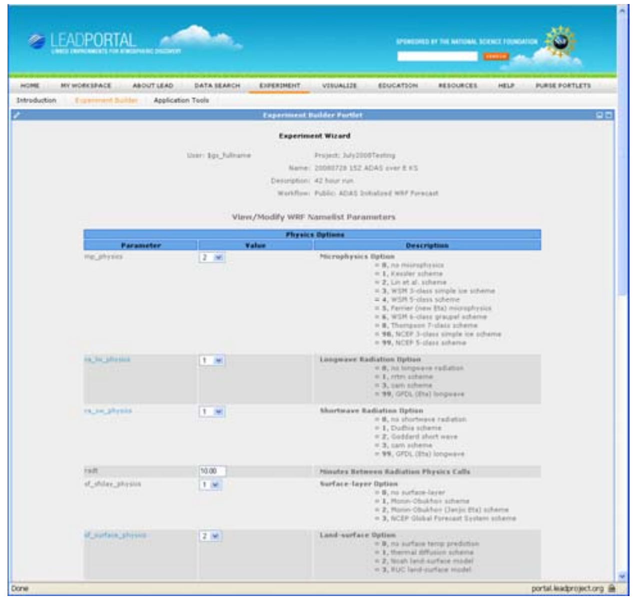

Figure 2.1. Experimental WRF namelist editor currently operational within the LEAD Portal.

One of the most important research components of the LEAD project involved developing and studying the behavior of workflows that are

- ***Event-driven*** (responding to weather or other triggers, especially real time CASA observations; Brotzge et al. 2006)
- ***Priority-based*** (execution urgency based upon user-assigned priority)
- ***Dynamically adaptive*** (change in response to the availability of resources, data)
- ***Fault-tolerant*** (ability to recover from application or resource failures while meeting the specified quality of service constraints)
- ***Data-driven*** (observations selected by the user and processed through ADAS or 3DVAR for analysis, visualization, or subsequent prediction using WRF)

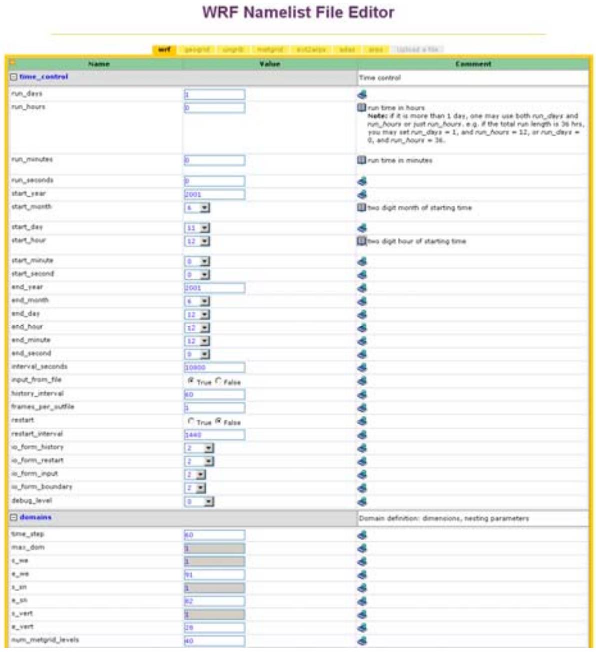

Figure 2.2. Experimental web-based namelist editor developed at CAPS for WRF, ADAS and other applications.

Two particular event-driven and dynamically adaptive workflows have been studied to date, one involving the mining of NEXRAD data with subsequent triggering of a WRF forecast based upon features identified, and the second involving the automated triggering of a WRF forecast based upon watches and mesoscale discussions (see §3). Fault tolerance already exists and was demonstrated at the LEAD Workshop held during the 2008 AMS Annual Meeting, and plans remain in place to bring the WRF 3DVAR package into LEAD, in addition to ADAS, for creating three-dimensional atmospheric analyses. Indeed, ADAS and the WRF 3DVAR system can be used sequentially (e.g., for single Doppler wind analysis), though such capability is beyond the scope of what can be accomplished during the no-cost extension. By adding the WRF 3DVAR system, we not only will enhance LEAD but make it much more attractive to the WRF community.

Related to data-driven workflows is the important task of creating a mechanism by which users can add their own data streams to the LEAD data system – in addition to the 17 data sets now available (Baltzer et al. 2008). This effort is ongoing in collaboration between Indiana University and Howard University.

Metadata and provenance of both data and scientific models is a necessary precondition for science reproducibility. LEAD supports this through a community XML schema for describing data objects and collections (i.e., the LEAD metadata schema: Plale et al. 2004, 2007; Simmhan et al. 2007), a service registry capturing attributes of models, and a metadata catalog (Jensen et al. 2008) that captures model configuration and results. The LEAD metadata schema has been enhanced and fully integrated into the LEAD framework, thereby allowing for fully automated metadata generation, including all products generated by workflows.

Sharing tools are being written to allow sharing of workflow metadata and provenance in other community metadata formats, as community needs dictate. The geospatial data query interface has been enhanced and users now can search their own personal workspace (myLEAD) for any products created or imported. In order to facilitate understanding, particularly by students, of the complexities of meteorological terminology, the Noesis ontology tool (Ramachandran et al. 2006) has been fully integrated as well. Additional data format and processing capabilities were added and enhancements made to facilitate crosswalks among various repositories and catalogs. In addition to many real time data sets now available within LEAD, the data system now contains online archives of certain products dating back several hundred days (Baltzer et al., 2008; Wilson et al., 2007, 2008).

The LEAD Portal interface (Gannon et el., 2007; Brewster et al. 2008; http://portal.leadproject.org) continues to be enhanced, and Java-based animation tutorials now are in place to lower the entry barrier, particularly for students. Additional enhancements have been added to facilitate greater flexibility in workflows, particularly with regard to parameter types. The workflow engine has been hardened and the NCSA Ensemble Broker (Alameda et al. 2007) has been linked to myLEAD, thus allowing users to generate hundreds of forecasts (ensembles or otherwise) and manage all of their results using LEAD data system tools. In fact, more than 1000 forecasts were so generated by NCSA during the spring, 2007 severe weather season (see Wilhelmson et al. 2008).

New monitoring and fault tolerance/recovery services have been added to LEAD. While results are encouraging, further layers of recovery are being added to **guarantee** that all workflows eventually complete successfully, and do so with no manual intervention. Several experiments have been conducted in which workflows are launched automatically based upon events detected in streaming data (Plale et al. 2007), and this has been made possible in part by advanced data mining and clustering techniques developed for application to NEXRAD Doppler radar data (Li et al., 2008). Research in dynamic adaptation using simple meteorological models, and sophisticated data assimilation systems, is providing guidance that will be used to determine how models might best be configured to produce an “optimal” forecast. This includes the ability to estimate forecast sensitivity and request additional observations so as to reduce the impact of observation errors. These capabilities will be used to help achieve a fully closed-loop dynamically adaptive capability.

3. APPLICATION OF LEAD IN THE 2008 NOAA HAZARDOUS WEATHER TEST BED

For the second year in a row, LEAD participated in the 2008 Spring Experiment of the NOAA Hazardous Weather Testbed in Norman, OK (Kain et al. 2008; see also http://hwt.nssl.noaa.gov/Spring_2008/), the principal goal of which is “to accelerate the transition of promising new meteorological insights and technologies into advances in forecasting and warning for hazardous mesoscale weather events throughout the United States.”

During the 2008 experiment, CAPS produced a 10-member real-time storm-scale ensemble with partial support from LEAD and partial support from a NOAA CSTAR grant (Xue et al 2008; Kong et al. 2008). At 4-km horizontal grid spacing, the WRF-ARW-based ensemble system ran daily during weekdays from 14 April through 6 June 2008 in a domain covering most of the continental U.S (Figure 3.1). Thirty-hour forecasts were produced starting at 00 UTC and extending through 06 UTC the next day. The ensemble system consisted of ten hybrid perturbation members composing perturbed initial conditions and various microphysics, shortwave radiation, and PBL physics parameterization scheme options. Additionally, a single 30-hour, 2-km deterministic forecast was produced over the same domain with radar data assimilation. All forecasts were run on the Cray XT3 at the Pittsburgh Supercomputing Center.

Several major changes from the 2007 experiment were made for the 2008 season, including: (1) the assimilation of NEXRAD Level-II data from over 120 radars via the ARPS 3DVAR and cloud analysis package into all but one ensemble member; (2) enlargement of the model domains (Figure 3.1); (3) initiation of daily 30 h forecasts at 0000 UTC using the NAM 12 km (grid 218) 00Z analyses as background for initialization, with the initial condition perturbations for the ensemble coming from the NCEP Short-Range Ensemble (SREF); (4) Use of both initial perturbations and physics variations in the ensembles.

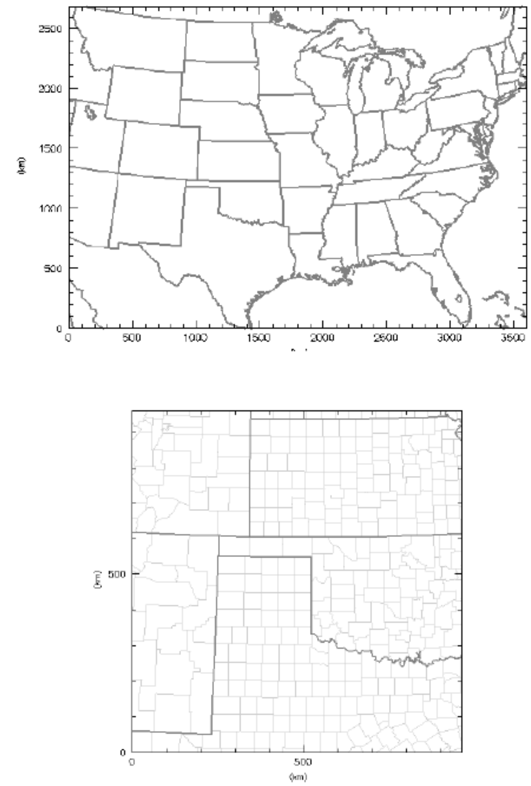

Figure 3.1. CAPS domain for 10-member, 4 km grid spacing WRF ensemble forecasts (top) and an example of an on-demand domain (bottom) for WRF forecasts run at 2 km grid spacing.

Ensemble forecast products were created in real time through existing capabilities in the SPC version of the N-AWIPS system for evaluation by researchers and operational forecasters during the experiment. During the course of the experiment, complete forecasts (i.e., all ensemble members plus the single deterministic run) were completed on 36 days. Compared to the 2007 season, the unprecedented assimilation of WSR-88D radar data nationwide indeed helped boost the equitable threat scores for the initial few hours with the overall impact of radar data lasting some 12 hours.

In addition to the ensembles and single high-resolution deterministic forecast coordinated by CAPS, NCSA and Indiana University, the LEAD on-demand effort ran concurrently from April 28 to June 6. Each day, a pair of forecasts was submitted over a domain of size 960 x 960 km at 2-km grid spacing. The initialization time was 15 UTC and forecasts were run for 15 hours, or until

06 UTC that evening (midnight CST). The objective was to measure the improvement in forecast skill using data from the morning of the severe weather forecast, compared to the ensembles that were initialized using data from 00 UTC from the evening before (6 pm CST). One member of the forecast pair was based upon the 3-hour forecast from the operational North American Mesoscale (NAM) model, the second from the LEAD operational 10-km ADAS analysis, the latter of which included Level II radar reflectivity data from the national NEXRAD network. The forecast domain was determined daily in consultation with Experiment participants, and the forecasts were submitted between 15 UTC and 16 UTC using the LEAD Portal. The Portal software controlled the run following manual domain selection, i.e., WRF model set-up and interpolation of initial and boundary data grids, job submission to the TeraGrid, and delivery of the results to via GridFTP. Once the files arrived, local software automatically created web pages and posted them to the CAPS web server for the forecasters to evaluate. On many days, the results, including the web page generating, were delivered in time for the 20 UTC final deadline for forecast submission and the 2015 UTC weather discussion.

The bulk of the 2008 LEAD on-demand WRF forecasts executed on TeraGrid resources at Indiana University, specifically on the BigRed supercomputer. Hardening of the file transfer methods during the summer and fall of 2007, in collaboration with the Globus and TeraGrid teams, and the use of BigRed for the execution, improved the robustness of the system over 2007 to the point where of the 62 forecasts submitted before 3 June 2008, 87% finished successfully. On June 4th a flash flood on the campus of the University of Indiana impacted the BigRed computer and also the LEAD data capacitor when power systems became unstable in the days following the flood. From June 4-6, five of six on-demand forecasts failed due to those problems.

The model runs are being verified against precipitation estimates from the 1-km NEXRAD precipitation mosaic prepared by the NOAA National Severe Storms Laboratory (Coniglio et al. 2008). Inter-comparisons are being performed and the on-demand forecasts also are being compared to statistics of the ensemble runs initialized at 00 UTC.

Prior to the hardening of data transfer software and other reliability enhancing efforts, the number of successful on-demand experiments made during the 2007 Spring Experiment was notably limited. During the past year, a subjective evaluation scoring system was created and 16 forecasts from the Spring 2007 campaign were scored subjectively. Preliminary results from this small sample are reported in Brewster, et al (2008).

Following the success of mining Storm Prediction Center (SPC) Mesoscale Discussions and Weather Watches via the SPC RSS feed to automatically launch on-demand nested WRF forecasts in 2007, a similar effort was undertaken in spring 2008 using the NCSA ensemble broker. Individual 27-km grid spacing forecasts were triggered based on SPC products, and the 2007 domain selection algorithm was modified slightly to eliminate duplication where the domain from a given proposed forecast might overlap significantly the domain of one previously made. A total of 476 WRF forecasts were submitted to the Abe TeraGrid supercomputer at NCSA during spring of 2008. In late June, following completion of the ensemble broker system, LEAD began testing an 8-member WRF ensemble based on mining real-time SPC products. The output for the triggered runs can be viewed at http://rt.atmos.uiuc.edu/trigger.

4. LINKAGES BETWEEN LEAD AND THE WRF PORTAL

a. Overview and Motivation

From its inception, LEAD has maintained close communication with the Joint NWP Test Bed (JNT), part of the Developmental Test Bed Center (DTC; http://www.dtcenter.org/) at NCAR, which is a national collaborative framework in which numerical weather analysis and prediction communities can interact to accelerate testing and development of new technologies as well as techniques for research applications and operational implementation – all in a way that mimics, but does not interfere with, actual forecast operations. In the early stages of LEAD, it became clear that the DTC was taking a different approach to modeling with a much shorter time horizon than LEAD, focusing on developing a desktop-client WRF configuration tool that executes WRF on a dedicated cluster.

As time progressed and both the LEAD and JNT efforts matured, conversations were restarted regarding possible strategies for integrating key elements from these two important projects. A more thorough technical analysis recently was conducted regarding integrating the JNT WRF configuration client into the LEAD infrastructure, and NSF has provided supplemental funding to support it. This fusion of portal technologies will allow JNT users to gain full benefit from the many aspects of LEAD that are not included in the WRF Portal (Govett and Smith 2008; see also http://www.wrfportal.org/index.html), -- such as access to grid computing resources, drag-and-drop assembly of complex workflows involving multiple codes, advanced, distributed data acquisition and management, access to dynamically adaptive models/sensors/tools, local data ingest, data assimilation, multi-level fault tolerance and recovery, data mining/feature extraction, ontology-based semantic search – while bringing to LEAD key capabilities related to namelist editing and WRF domain configuration, including (multiply) nested grids.

b. Goals

The foremost goal of the LEAD-JNT integration project is to bring the JNT's WRF Domain Wizard (WDW) configuration client (Smith et al. 2008; see also http://www.wrfportal.org/DomainWizard.html) into the LEAD infrastructure. This will provide LEAD users with the ability to edit namelist input files and incorporate them into workflows for key applications such as WRF, WPS, WRF 3D-VAR, and ADAS. The enhanced capabilities of the WDW will greatly expand the ways in which key applications can be used in LEAD, ranging, in the case of WRF, from simple idealized configurations (e.g., a synoptic-scale baroclinic wave channel model, flow over a bell-shaped mountain, a 2D gravity current, a cloud model, or a dry boundary-layer model) to the very sophisticated (multiply nested vortex-tracking moving grid simulations of hurricanes, parametric studies involving the assimilation of multiple types of observations).

The second major goal is to develop a new namelist input file error-checking/validation tool, infused with numerical modeling intelligence to prevent the generation of illogical, corrupt namelist files, and incorporate this capability into LEAD workflows. The tool will guide users through parameter selection, drastically reduce user frustration and conserve computing resources by ensuring that simulations submitted to remote job queues don't fail due to easily preventable errors (CFL stability violations, excessively small MPI patch size, specification of incompatible physics options, incorrect grid nesting hierarchies and layouts, inconsistent dates, etc.).

The third goal is to enhance the WRF Portal's vertical editor to include options for stretching of vertical coordinate surfaces. These configuration options have proven invaluable for optimizing WRF simulations, making possible a reduction in the total number of vertical levels while retaining ample resolution in dynamically active regions. The precision of the WRF Portal's horizontal editor and the quality of the graphics it renders also will be enhanced to improve the centering of multiply nested grids.

Closely related is the capability for users to edit, compile and manage their own versions of application (e.g., WRF) source codes – a capability which now exists but has not yet been exposed to users through the LEAD Portal. Also high on the list of priorities is a closer linkage from the LEAD portal to the NCSA ensemble broker system to support "parametric workflows," in which selected parameters of the WRF can be varied systematically over a specified range, or randomly within an ensemble envelope, in a single workflow (e.g., testing various values of computational mixing or user-specified parameters in physics parameterizations).

A summary of the capabilities being transferred between the LEAD and DTC portals is presented in Figure 4.1 and list of new capabilities being developed for both portals is shown in Figure 4.2.

c. Status

Shown below is the status of key tasks associated with the LEAD-WDW integration effort.

- *Add WRF Domain Wizard Java Web Start launch capability to LEAD Portal.* LEAD and NOAA scientists felt that the most effective initial approach for linking the LEAD and WRF portals was to "loosely couple" them, and this is done via a single WDW JAR (Java archive) file on the WRF Portal website and two different JNLP files that launch different versions of the application. NOAA/ESRL has created a separate website to serve the WRF Domain Wizard for LEAD, and Figure 4.3 shows WDW-LEAD running independently of both WPS and a remote cluster connection. Note that the namelist editor now includes a new "Validate" button that runs a series of basic error-checking functions. An "Errors Found" dialog box displays a list of all errors detected in a user's WRF namelist.input file. Each item includes the parameter name, its incorrectly specified value, an error message that provides guidance on how to correct the error, and the line number at which the error occurs. Efforts are underway to install a hyperlink in the myLEAD workspace to access the new WDW-LEAD web site.

Capabilities Transferred from LEAD → WRF Portal	Capabilities Transferred from WRF Portal →LEAD
Web-based namelist input file editor	Desktop namelist input file editor
Ability to submit large simulations across the TeraGrid, access to NSF computing resources	Task/script/runtime environmental variable editors
Drag-and-drop workflow assembly	Ability to run parametric workflows
Distributed data acquisition and management	`diff` tool for comparing/summarizing workflows
Access to dynamically adaptive models, sensors, and tools	WRF nesting: multiple, two-way interactive, vortex-tracking/moving nests
Advanced data assimilation (ADAS, WRF 3D-VAR, GSI)	Elegant desktop GUI with rich interactive graphics environment
Multi-level fault tolerance and recovery	
Data mining and feature extraction	
Ontology-based semantic search	
Improvements to logic and precision of the WRF Portal's horizontal editor to improve the centering of multiply-nested grids	

Figure 4.1. Capabilities being transferred between LEAD and the WRF Portal.

New Portal Capabilities
Namelist file error-checking/validation tool to prevent input/configuration errors
Edit, compile and manage application source code (expose to user)
Vertical grid stretching tool
Artificial robotic test-subject with multiple personalities for system evaluation

Figure 4.2. New capabilities being developed in both the LEAD and WRF Portals.

- *Develop automated parameter validation capabilities for the WRF model namelist input files*. In order to make rapid progress on the development of a stability checker for WRF namelist.input files, it was decided to initially employ practical "rules of thumb" rather than implement complicated CFL criteria based on the numerical schemes selected, their accuracy, the numerical order of the computational viscosity chosen, etc. For the WRF-NMM core, the rule of thumb is to set the time step (in seconds) to approximately 2.25 times the horizontal grid spacing (in km). Thus, if the horizontal spacing of a grid domain is 12 km, then the time step should be 27 seconds or less. For the WRF-ARW core, the rule of thumb is to set the time step to 4-6 times the horizontal resolution (in km), so 72 seconds would be the recommended time step for the aforementioned grid domain. Both of these rules were coded into the new "isStable" function. A new "integralTimeStep" function also was developed to check whether the user specified value of Δt will result in an integral number of time steps per hour. If not, this function suggests a new value that will satisfy the criterion.

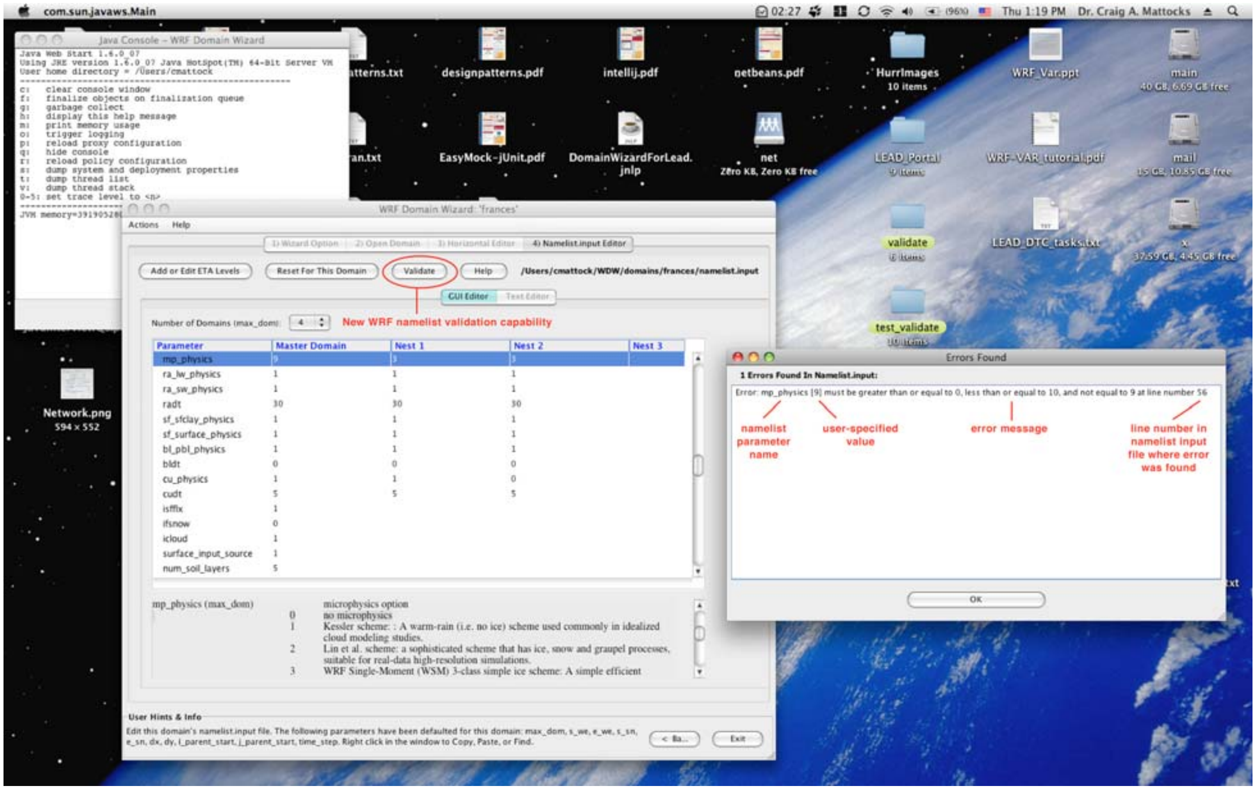

Figure 4.3. Screenshot of the LEAD version of the WRF Domain Wizard running independently of both the WRF Processing System and an SSH connection to a remote supercomputing cluster. The "Validate" button (red circle) allows users to query namelist input files for configuration errors, as indicated by the "Errors Found" window on the right portion of the figure.

Other tasks, in addition to enhancing the vertical grid editor, include developing automated parameter validation capabilities for nested grids, model physics, and external data sets.

5. PLANS FOR COMMUNITY DEPLOYMENT

From the beginning, the LEAD vision has been to not only conduct excellent research and develop useful technologies, but to do so in a practicable way that transforms meteorological research and education and has value to other disciplines. However, the transformation can occur only when LEAD is transitioned from a research project and made available as a persistent, sustained facility upon which the community can rely. This notion was expressly stated in the original LEAD proposal:

> *"As a virtual extension of the user's desktop, LEAD will enable researchers, users, educators, and students to use atmospheric models and other tools in more realistic, real time settings than is now possible."*

an effort to enhance its visibility to the community, among other things, the LEAD team presented a 90-minute tutorial as part of the annual WRF Workshop in Boulder, CO on June 27, 2008. It was attended by some 30 participants, most of whom brought a laptop and were able to access the LEAD Portal, construct and submit WRF based workflows and view results. An important capability made available by LEAD was namelist editing for the WRF model. Further, this was the first time LEAD had a presence at the WRF Workshop, and it represented an opportunity to explain how LEAD and the WRF Portal differ, and how both groups are working to leverage and couple their capabilities.

LEAD has been used in a variety of education initiatives, including WxChallenge (Clark et al. 2008), and options are being explored for its application by campus weather services at colleges and universities. Other potential roles for LEAD include on-demand, tailored numerical forecasting for guiding crew and instrument deployment in field programs (discussions now are underway with leaders of VORTEX2), and it LEAD already is being used in crop modeling research by Western Michigan University Scientists as part of a project funded by USDA.

The outcomes and impacts realized by LEAD, only a few of which were mentioned herein, provide a strong foundation upon which to build a persistent cyberinfrastructure for atmospheric sciences, with extensibility to other science and engineering domains. Consequently, funding is being sought to continue making LEAD available to the community through the end of calendar year 2009, and longer-term plans for deployment also are being developed.

6. ACKNOWLEDGMENTS

LEAD is 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), ATM-0331586 (Howard University), ATM-0331587 (UCAR), and ATM-0331578 (University of Illinois at Urbana-Champaign with a sub-contract to the University of North Carolina). CASA is funded in part by the Engineering Research Centers Program of the National Science Foundation under NSF Cooperative Agreement EEC-0313747 to the University of Massachusetts-Amherst. Support for the NOAA Hazardous Weather Test Bed was provided by NOAA C-STAR Grant NA17RJ1227/NAGO000070029 to the University of Oklahoma. We gratefully acknowledge Drs. Mark Govett and Jeff Smith, NOAA/ESRL, for their enthusiasm for and substantial assistance with the LEAD-DTC integration effort.

7. REFERENCES

Alameda, J. and Co-Authors, 2007: Siege: A graphical user interface to enable management of large numbers of weather simulations. Preprints, *23rd Conf. On IIPS*, 15-18 January, San Antonio, TX, Amer. Meteor. Soc.

Baltzer, T. and Co-Authors, 2007: LEAD at the Unidata workshop: Demonstrating the democratization of NWP capabilities. Preprints, *23rd Conf. On Integrated Information and Processing*, 15-18 January, San Antonio, TX, Amer. Meteor. Soc.

Baltzer, T., and Co-Authors, 2008: The LEAD test bed system at the Unidata Program Center: a medium term archive of meteorological data. Preprints, *24th Conf. On Integrated Information and Processing*, 21-24 January, New Orleans, LA, Amer. Meteor. Soc.

Brewster, K. and Co-Authors, 2008: Use of the LEAD portal for on-demand severe weather prediction. Preprints, *24th Conf. On Integrated Information and Processing*, 21-24 January, New Orleans, LA, Amer. Meteor. Soc.

Brotzge, J., K.K. Droegemeier, and D.J. McLaughlin, 2006: Collaborative Adaptive Sensing of the Atmosphere (CASA): New radar system for improving analysis and forecasting of surface weather conditions. *J. Transport. Res. Board*, No. 1948, 145-151.

Burgess, D. W., D. C. Dowell, L. J. Wicker, A. Witt, 2005: Detailed comparison of observed and modeled tornadogenesis. Preprints, 32nd Conf. on Radar Meteor., Albuquerque, NM, Amer. Meteor. Soc.

Clark, R. D., S. Marru, M. Christie, D. Gannon, B. Illston, K. Droegemeier, and T. Baltzer, 2008: The LEAD-WxChallenge Pilot Project: Enabling the Community. (Preprints). 24th International Conference on Interactive Information Processing Systems. American Meteorological Society, New Orleans, LA, January.

Coniglio, M.C. and Co-Authors, 2008: Evaluation of WRF model output for severe-weather forecasting from the 2008 NOAA Hazardous Weather Testbed Spring Experiment. Preprints, 24th Conf. on Severe Local Storms, Savannah. GA, Amer. Meteor. Soc., Paper 12.3.

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, 2007: A new paradigm for mesoscale meteorology: Grid and web service-oriented research and education in LEAD. Preprints, *23rd Conf. On Integrated Information and Processing*, 15-18 January, San Antonio, TX, Amer. Meteor. Soc.

Droegemeier, K.K., 2008: Transforming the Sensing and Numerical Prediction of High Impact Local Weather Through Dynamic Adaptation. *Phil. Trans. of the Royal Soc*. A, 1-20.

Gannon, D., and Co-Authors, 2007: The LEAD science portal problem solving environment. Preprints, 23rd Conf. On Integrated Information and Processing, 15-18 January, San Antonio, TX, Amer. Meteor. Soc.

Govett, M.W. and J.S. Smith, 2008: A public release of the WRF Portal: A graphical user interface for WRF. Preprints, *24th Conf. On Integrated Information and Processing*, 21-24 January, New Orleans, LA, Amer. Meteor. Soc.

Jensen, S. and B. Plale 2008: Using Characteristics of Computational Science Schemas for Workflow Metadata Management, IEEE 2008 Second International Workshop on Scientific Workflows (SWF'08), Honolulu Hawaii.

Kain, J.S., S.J. Weiss, D.R. Bright, M.E. Baldwin, J.J. Levit, G.W. Carbin, C.S. Schwartz, M. Weisman, K. Droegemeier, D. Weber, and K.W. Thomas, 2008: Some practical considerations for the first generation of operational convection-allowing NWP: How much resolution is enough? *Wea. and Forecasting*, **23**, 931-952.

Kong, F., M. Xue, K.W. Thomas, K.K. Droegemeier, Y. Wang, K. Brewster, J. Gao, J. Kain, S.J. Weiss, D. Bright, M. Coniglio, and J. Du, 2008: Real-time storm-scale ensemble forecast experiment: Analysis of spring 2008 experiment data. Preprints, 24th Conf. on Severe Local Storms, Savannah. GA, Amer. Meteor. Soc., Paper 12.3.

Kuhlman, K.M., T.M. Smith, G.J. Stumpf, K.L. Ortega, and K.L. Manross, 2008: Experimental probabilistic hazard information in practice: Results from the 2008 EWP spring program. Preprints, 24th Conf. on Severe Local Storms, Savannah. GA, Amer. Meteor. Soc., Paper 12.2.

Li, X., and Co-Authors, 2008: Storm clustering for data-driven weather forecasting. Preprints, *24th Conf. On Integrated Information and Processing*, 21-24 January, New Orleans, LA, 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

Plale, B., D. Gannon, J. Brotzge, K.K. Droegemeier and Co-Authors, 2006: CASA and LEAD: Adaptive cyberinfrastructure for real-time multiscale weather forecasting. *IEEE Computer*, **39**, 66-74.

Plale, B. and Co-Authors, 2007: Real time filtering and mining of NEXRAD streams for mesoscale forecast and prediction. Preprints, *23rd Conf. On Integrated Information and Processing*, 15-18 January, San Antonio, TX, Amer. Meteor. Soc.

Ramachandran, R., S. Movva, P. Cherukuri, and S. Graves, 2006: Noesis: An Ontology-based Semantic Search Tool and Resource Aggregator. Geoinformatics Conference 2006, Reston, Virginia.

Simmhan, Y.L., Sangmi Lee Pallickara, Nithya N. Vijayakumar, and Beth Plale, 2007: Data Management in Dynamic Environment-driven Computational Science, *IFIP Federation for Information Environments, Vol. 239. Grid-based problem solving environments.* P.W. Gaffney and J.C.T. Pool, Eds, Springer Boston, pp. 317-333.

Smith, J. S., P. McCaslin, and M. W. Govett, 2008: WRF Domain Wizard: The WRF preprocessing system GUI, 88th AMS Annual Meeting, New Orleans, La., Amer. Meteor. Soc., 21 pp.

Stumpf, G.J., T.M. Smith and K.A. Scharfenberg, 2007: The future of US severe weather warning operations. Preprints, 4th European Conf. on Severe Storms, 10-14 September, Trieste, Italy, 2 pp. (Available at http://www.essl.org/ECSS/2007/abs/06-Forecasts/1177965431.stumpf-1-sec06.oral.pdf.)

Weiss, S. J., J. S. Kain, D. R. Bright, J. J. Levit, G. W. Carbin, M. E. Pyle, Z. I. Janjic, B. S. Ferrier, J. Du, M. L. Weisman, and M. Xue, 2007: The NOAA Hazardous Weather Testbed: Collaborative testing of ensemble and convection-allowing WRF models and subsequent transfer to operations at the Storm Prediction Center. *22nd Conf. Wea. Anal. Forecasting/18th Conf. Num. Wea. Pred.*, Salt Lake City, Utah, Amer. Meteor. Soc., CDROM 6B.4.

Wilhelmson, R.B.. and Co-Authors, 2008: Automatic triggering of high-resolution forecasts in response to severe weather indicators from the NOAA Storm Prediction Center. Preprints, *24th Conf. On Integrated Information and Processing*, 21-24 January, New Orleans, LA, Amer. Meteor. Soc.

Wilson, A., T. Baltzer and J. Caron, 2007: The THREDDS data repository (TDR) for long term data storage and access. Preprints, *23rd Conf. On Integrated Information and Processing*, 15-18 January, San Antonio, TX, Amer. Meteor. Soc.

Wilson, A., and Co-Authors, 2008: Create an archive with the THREDDS data repository. Preprints, *24th Conf. On Integrated Information and Processing*, 21-24 January, New Orleans, LA, Amer. Meteor. Soc.

Xue, M., F. Kong, K.W. Thomas, J. Gao, Y. Wang, K. Brewster, K.K. Droegemeier, J. Kain, S. Weiss, D. Bright, M. Coniglio, and J. Du, 2008: CAPS realtime storm-scale ensemble and high-resolution forecasts as part of the NOAA Hazardous Weather Testbed 2008. Preprints, 24th Conf. on Severe Local Storms, Savannah. GA.