Unidata equipment grant - University of Miami

A proposal submitted to Unidata Program Center (UPC) Equipment Awards

Principal Investigator:
Dr. Brian Mapes, Associate Professor
Meteorology and Physical Oceanography
University of Miami - Rosenstiel School of Marine and Atmospheric Science
4600 Rickenbacker Causeway, Miami, FL 33149-1031
phone: (305) 421-4275
fax: (305) 421-4696
email: mapes@miami.edu

Co-PI:
Craig Mattocks (University of Miami)

Project Dates:
May 1, 2013 – April 30, 2014

Budget:
University of Miami: $18,000

Institutional Representative:
Fernande Saintilis
Sponsored Programs, Team Manager
University of Miami - Rosenstiel School of Marine and Atmospheric Science
4600 Rickenbacker Causeway, SA 111
Miami, FL 33149-1031
phone: (305) 421-4181
fax: (305) 421-4876
email: fsaintilis@miami.edu
DUNS: 152764007

Signature: _____ **Date:** 3-13-13

A. Project Summary:

We propose installing Unidata's full technology suite at the University of Miami, including both of the Special Consideration items detailed in the RFP (http://www.unidata.ucar.edu/community/equipaward/RFP2013.html).

The AWIPS-II system will be used for education (classroom and graduate) and research (including hurricane work and field measurements support). But it will also provide a novel pathway for research to operations (R_2O) collaborations with the National Hurricane Center (NHC), through Java/Javascript code development to create new AWIPS-II products and services that are accessible across multiple enterprises. A data server will take IDD/ LDM ingests and put out TDS and EDEX services, including server-side Web graphics production via RAMADDA/IDV, and also will host a diverse set of locally generated and offered data.

Data served will include at least a 3-month archive of all standard operational weather products including radar. In addition, we will host climate and weather research datasets, such as model outputs, and locally generated field data (including lidar, radar, webcams, and time series) from Miami's NSF-funded Cloud-Aerosol-Rain Observatory (CAROb). A lively and growing community of IDV users on site (both students+instructors and researchers) will develop and post catalogs of data-integration bundles and case studies that attempt to span the weather-climate gap, as well as the gap between operational meteorology and local special observations. All bundles and data will be shared and advertized with the goal of being a best-practices model for reproducibility and openness in atmospheric science. The PI will take Unidata's summer trainings July 24 - August 8, 2013 and will be in Boulder for the summer to interview Unidata personnel as needed, to ensure sufficient local knowledge for a successful Miami implementation of Unidata's complete technology suite. Several other interested faculty will ensure sufficient depth of interest in the server to sustain it over time.

B. Project Description

The University of Miami's Division of Meteorology and Physical Oceanography seeks to jump to the forefront of Unidata's technologies with a new central server computer and full suite of software. Although initially driven by the interests of this PI (Prof. Brian Mapes) in redesigning the teaching of our weather-related classes ("Weather Analysis" for undergraduates, "Convective and Mesoscale Meteorology" for graduate students, and others taught by other faculty), our impetus to jump ourselves to the forefront of Unidata technologies now has multiple other sources:

- Longtime users of McIDAS, often for hurricane field campaign forecasting and nowcasting (groups of Profs. Shuyi Chen and Sharan Majumdar), are seeking to update their capabilities (adding more IDV efforts for example, and likely moving to AWIPS II as it proves out).
- A skilled Java/Javascript developer here (Dr. Craig Mattocks) seeks pathways to transition software advances to NHC. He is currently collaborating with NHC and Prof. Beth Plale, Director of Indiana University's Data-to-Insight Center, on a NSF-funded project to get the SLOSH storm surge prediction model running on cloud-computing (CC) resources. This project has reached a stage in its development where the forecasters in NHC's Storm Surge Unit now want to configure and submit ensemble simulations through a GUI to a remote CC service such as Amazon's Elastic Cloud (EC2, http://aws.amazon.com/ec2/), then composite and visualize the output to create maximum envelopes of water (MEOWs/MOMs) and high-resolution inundation maps. Since AWIPS II will become the de facto operational hub for accessing gridded model data throughout NOAA's national centers and NWS forecast offices, it would be ideal if the front end of this CC system could be developed for AWIPS II. With its new Java-based Eclipse plug-in and Javascript-based XML/XUL capabilities, the AWIPS Development Environment (ADE) offers proficient solutions for GUI tool development. The CC workflow pipeline middleware software developed in the NSF project could be contributed to enhance the computational power of the AWIPS II system. Of course all such improvements would also be shared broadly with the entire Unidata community.

- High-resolution reference "Nature Run" WRF output datasets from Prof. David Nolan, produced for a community-based OSSE (Observing System Simulation Experiment) project, need to be served out to collaborators, as well as visualized in common views that the whole community can see in the same way. A RAMADDA server with the data and relevant IDV bundles seems an excellent solution, as borne out by pilot attempts on our test RAMADDA.
- The Cloud-Aerosol-Rain Observatory (CAROb) of Profs. Paquita Zuidema and Bruce Albrecht, a major instrumentation project from NSF, seeks ready outlets for its datasets and their visualizations. These local observations need to be put in the context of the weather, and again a RAMADDA with IDV bundles solution, with ISL-generated visualization thumbnails on a realtime web page, seems ideal.
- Several UM climate and modeling groups, including that of PI Mapes, share their data internally in cumbersome ways. All can easily be persuaded that a Unidata solution is superior. Some efforts include data-assimilative work that should be integrated in IDV bundles with verification data (differences etc.). These users will help advance statistical dimensions of the IDV, such as climate uses of the ensemble functionality, which may help all IDV users.
- The PI (Mapes) is attempting to assemble integrated visualizations for past cases of extreme events worldwide, using the diverse and growing set of OpenDAP global datasets online in TDS and GDS servers, for both research and pro bono WMO-based developing-world forecaster training activities. I anticipate submitting a major NSF proposal within the year to study extreme events in the context of climate probability, building tools to allow users to drill down from histograms to fully detailed synoptic style analyses with little or no user programming. A well-curated catalog of easily-adjusted richly detailed bundles served from a RAMADDA page is my intended technological base for the project. These tools will benefit the whole Unidata community.
- UM is proposing to host the NCAR Advanced Study Program (ASP) Summer Colloquium on Hurricanes and their Impact, in Summer 2014. The new capabilities will potentially form a central part of the lectures and lab exercises on hurricane

observations, modeling, ensemble prediction, diagnostics and forecasting for the 25-30 students attending the Colloquium.

Oceanographers in our Division have also expressed interest, but they know too little to join this proposal meaningfully at this stage. With successful implementation, however, we anticipate a significant interest arising in this allied field. If this proposal is supported, the capabilities will be the subject of a Divisional seminar in autumn 2013 by the PI (an enthusiast who will be fresh from the Unidata August 2013 training) to try and spark the transitional efforts.

Beyond our walls, these datasets and services and our experiences transitioning to them will be shared freely and advertised proudly to the communities with whom we interact. We plan to arrange a regional training in Miami in about a year, and have received positive indications from Unidata staff about that prospect. While we are primarily a research and teaching institution rather than a service-providing one, the open quality of Unidata's approach will naturally benefit others. This is good recruiting and public relations, and we do not anticipate traffic being great enough to engender any conflicts between our School's missions and objectives and the resulting, related free services traffic.

The equipment requested, based on advice from Tom at Unidata-Egrants support, is for 3 servers. One will be primarily a data server, with hardware choices modeled on the new backup for motherlode at UCAR to minimize any new setup difficulties. Quoting tom@Unidata again, "we are hoping that the machine(s) capable of decoding all data flowing in the Unidata IDD (includes all data broadcast in the NOAAPort SBN and other feeds like CONDUIT (high resolution model output from NCEP) and NEXRAD Level II volume scan data) will be in the $5K-10K range." This will replace an aging server (weather.rsmas.miami.edu), which has been our LDM server for years but is poorly indexed and of little use except as a sandbox for playing with our first RAMADDA instance (installed in December 2012).

The other two requested machines will be dedicated more to AWIPS-II use and to power-user experimentation and development. For these purposes they will likely with configured to emphasize more processing and graphics. Again our choices will be based on

Unidata's successful configuration as reported by Tom and detailed in the Budget section below.

The most invaluable aspect of this grant from our point of view will be Unidata's technical expertise in setup and installation, so we are very open to choosing whatever hardware supports easy installation and low future maintenance.

Existing computing resources here are numerous, with several Linux based clusters owned by various research PIs hosted in a central server room of the building. As a result, power and networking and rack space and basic local expertise are conveniently available.

Personnel, expertise and institutional setting

PI Mapes is primarily an educator and researcher, concerned foremost with user experiences and capabilities, but with many years of user-level experience and skills in computing and scientific programming. He serves on the IDV steering committee and is a frequent contributor of questions to the support-idv help ticket archives. He will drive the project forward with a desire for final capabilities and products, gaining and entraining expertise as needed.

Co-PI Mattocks has over 30 years of experience as a numerical modeler and software developer. He has served as PI, Co-PI, and Manager in the design and development of several operational numerical weather/ocean prediction systems. He is particularly adept at developing and implementing advanced algorithms in numerical models. A tenacious, task-oriented problem solver versed in extreme/agile programming and software architecture, he thrives on producing clean, high-quality, well-documented computer code in both procedural and object-oriented languages, including Fortran, C/C++, Java, Javascript, and HTML/CSS. He has developed custom Java GUI components and dynamic HTML animation capabilities for applications and web sites in the private sector.

The Rosenstiel School of Marine and Atmospheric Sciences (RSMAS) has a faculty of about 80, and total research computer users in the hundreds. The RSMAS Computing Facility (RCF) has been a semi-autonomous unit of the University of Miami, with University-leading skills due to the data-intensive nature of our field. Closer integration with the University is occurring now, which will stabilize RCFs budget and mission in the

years ahead. In brief we believe the facilities (server room) and staff will be sufficient for a successful outcome of this proposal.

C. Budget

The total requested funding for this proposal is $18,000, which as capital equipment is exempt from indirect costs. This will allow for the purchase of 3 rack mounted systems as discussed above: Based on matching or exceeding the capabilities of motherlode.ucar.edu, we seek

1. One dual quad core headless box (no video card) machine that has >64 GB of RAM and at least 20 TB of mirrored disk
2. Two powerful graphics-capable machines, one for anytime AWIPS II usage, one for development and experimentation. Based on discussions with Tom at the Unidata Program Center and some initial vendor quotes, we are proposing servers with dual quad core CPUs, 24 GB RAM or more, Nvidia graphics, and >1 TB storage split across 4 disks with stripped RAID, running 32-bit RedHat Enterprise Linux.

By removing line items from our preliminary hardware vendor quotes (academic pricing, no sales tax), it appears that Machine 1. can be obtained for about $8220, while the servers under point 2. can be obtained for $5324 each, yielding a total of $18,868 that falls safely under the program limit. Server-room rack space, electricity and cooling will be provided by the School as an expense of its research and instructional mission. The PI and CoPI will contribute their own time and travel costs for attending Unidata Software Training Workshops, managing the Unidata hardware and software, and transferring this knowledge to faculty, students and other interested members of the UM-RSMAS community. This time is considered an investment in future teaching quality and in facilitating proposals under development as described above.

D. Project Milestones

We expect a timeline like the following:

• May 2013 – Notification of award.

- July 2013 – Acquire and install equipment.

PI attends July-August Unidata trainings.

- August 2013 – Set up and configure LDM, CAVE/EDEX/AWIPS, THREDDS/TDS/IDD, and RAMADDA software systems.
- Fall 2013 – System fully operational, PI teaches Weather Analysis using it.

- October - November 2013 – PIs and graduate students attend fall Unidata Software Training Workshops in Boulder, CO to update their skill sets with Unidata software, learn how to administer/manage both the hardware and software, and transfer this knowledge to the UM-RSMAS community.
- December 2013 – Start serving high-resolution WRF hurricane Nature Run simulation results with collaborators at NOAA/AOML's Hurricane Research Division, NHC, and the broader university research community through RAMADDA. Provide IDV bundles to ease visualization of this extremely rich dense dataset.
- January 2014 – Begin development to create custom GUI to enable NHC Storm Surge Specialists to submit SLOSH storm surge predictions on remote cloud resources from AWIPS II, then composite and visualize the results.
- Winter-Spring 2013-14: Transfer Hurricane nature runs, CARCOb observatory data, climate model ensembles, and other local datasets to RAMADDA server. Develop thumbnail and showcase IDV bundles as index material for these data resources.
- April 2014 – Submit a brief final report, in the form of a short article, to the Unidata Program Center that describes: 1) how the computer hardware was employed, 2) the key capabilities that Unidata software provided, 3) the benefits of participating in the Unidata Equipment Awards program, 4) how working closely with Unidata accelerated the development of new capabilities and their transition to an operational forecasting environment, 5) the impact of delivering unique ultra-high resolution numerical model datasets to the educational, research, emergency management, and ecosystem restoration communities, and 6) assessments on approaches that worked/failed in order to improve the program.