DocuSign Envelope ID: 44B28206-6199-4D80-9DAB-4E549E8B166C

New proposal to the Unidata Program Center

for

Upgrading THREDDS and Deploying JupyterHub at the University of Wisconsin-Milwaukee to Support Education and Research

by

The Board of Regents of the University of Wisconsin System
on behalf of the University of Wisconsin-Milwaukee
P. O. Box 340
Milwaukee, Wisconsin 53201-0340

PRINCIPAL INVESTIGATOR: A. Clark Evans
(with support from graduate students Alex Moxon and Michael Vossen)

PERIOD OF ACTIVITY: 1 June 2021 – 31 May 2022

SUBMITTING DATE: 22 March 2021

AMOUNT REQUESTED: $10,672

PRINCIPAL INVESTIGATORS AND ENDORSING OFFICIALS:

A. Clark Evans (Lead PI)
Associate Professor
University of Wisconsin-Milwaukee
(414) 229-4469, evans36@uwm.edu

DocuSigned by:
Kate Mollen
555DCF8C12B6453...

Kate Mollen (UWM Endorsing Official)
Director, Office of Sponsored Programs
University of Wisconsin-Milwaukee
(414) 229-4853, grant-notice@uwm.edu

B. Project Summary

The University of Wisconsin-Milwaukee (UWM) Atmospheric Science Program, a member of the University Corporation for Atmospheric Research since 1979, offers B.S., M.S., and Ph.D. degrees in atmospheric science. We also offer the unique cocurricular Innovative Weather (Roebber et al. 2010) forecasting and decision-support program and robust undergraduate research opportunities alongside our faculty. Approximately 35% of our program's graduates enter the National Weather Service, 35% pursue other public-sector employment (e.g., in academia, with federal research labs, or at state agencies), and 30% pursue private-sector employment (e.g., as broadcast meteorologists, risk assessors, and forecasters) following graduation.

Our program's current computing resources include a small desktop computer lab used by selected courses and for non-data-intensive research; a Dell PowerEdge T710 server purchased in 2012 that runs LDM, THREDDS, and RAMADDA; a Dell PowerEdge T430 server and three Dell Precision T1700 workstations purchased in 2015 used as an EDEX server and CAVE clients, respectively; and the campus *mortimer* supercomputer for data-intensive research. Unfortunately, the nine-year-old Dell PowerEdge T710 server has experienced multiple hardware failures in the last two years, and our other resources are not sufficiently equipped to handle and/or explicitly do not support the activities for which this server is used. Furthermore, coincident with larger programmatic changes, we desire to deploy a local JupyterHub instance in support of implementing modern Python-based scientific workflows throughout our educational curricula. Again, however, our existing resources are not sufficiently equipped to handle and/or explicitly do not support a local JupyterHub instance.

As a result, to enable us to continue our contributions to the Unidata community and improve the educational resources available to our students, we seek assistance from Unidata to purchase a new Dell PowerEdge T640 server. This machine will allow us to continue our IDD participation, share data with the community largely through THREDDS, and deploy a JupyterHub instance to support data-driven education in multiple undergraduate and graduate courses and our Innovative Weather cocurricular student training program. We request Unidata funding since our department, college, and university unfortunately have not had, do not have, and do not expect to soon have the needed funding to support this purchase.

C. Project Description

a. Introduction

The department and institutional computing resources currently available to UWM's Atmospheric Science program faculty and students are as follows:

- Twelve iMacs purchased in 2014. These machines support multiple courses (e.g., Synoptic Meteorology I and II, Mesoscale Meteorology, Climate Dynamics, Statistical Methods in the Atmospheric Sciences I and II, Data Analytics, and Air Pollution Modeling) and non-resource-intensive student research.
- A Dell PowerEdge T710 server purchased in 2012 with funding from a Unidata Equipment Award. This server allows us to participate in Unidata's IDD program and share data with the community using Unidata and community technologies (THREDDS and RAMADDA).

This server is also used as a student training platform by our Innovative Weather program, wherein students have gained expertise in generating meteorological analyses from model data and observations (mostly using GEMPAK, but recently also using Python) in support of Innovative Weather's forecast and risk-assessment operations.

- A Dell PowerEdge T430 server and three Dell Precision T1700 workstations purchased in 2015 with funding from a Unidata Equipment Award. The server is used as an EDEX server and the workstations are used as CAVE clients to facilitate student training on AWIPS-II, as reported on in a May 2016 News@Unidata feature.
- The *mortimer* supercomputer, operational since 2015, comprised of 2,724 compute cores, 11.3 TB RAM, and 738 TB RAID storage. *Mortimer* is used for resource-intensive research by faculty and graduate students.

This proposal is motivated by the need to upgrade the nine-year-old Dell PowerEdge T710 server to maintain our IDD participation, continue to share locally produced datasets with the community, and to better support the integration of Python-based modern scientific workflows throughout our undergraduate and graduate curricula. To the former two points, though the PowerEdge T710 has served our program well for nine years, it has experienced multiple hard drive failures over the last two years and will likely experience further component failures as it continues to age. This imperils our ability to continue to participate in the IDD and share data with the community via THREDDS and RAMADDA.

To the latter point, consistent with the move toward Python to fulfill the programming requirement in many atmospheric science undergraduate curricula, our program is transitioning to Python from FORTRAN for its programming requirement effective in fall 2021. Coincident with this transition, we wish to modernize the meteorological analysis components of several upper-level and graduate courses. However, as many of our students (particularly at the undergraduate level) work full-time to support themselves through college, they do not always have access to on-campus resources to complete their coursework, nor do they necessarily have sufficiently robust technology at home to support data-intensive coursework. A local JupyterHub instance can address these issues; however, our existing servers are not sufficiently powerful to support several concurrent users connected to a local JupyterHub instance at the same time as their other processes (e.g., LDM, THREDDS, and EDEX), and there is no local funding to support purchasing a server to facilitate these activities.

Consequently, to enable us to continue our IDD participation, share data with the community using THREDDS, and better train our students in modern Python-based scientific workflows, we request Unidata funding to purchase a Dell PowerEdge T640 server. This server will include 2x Intel Xeon Gold 5218 processors running at 2.3 GHz with 32 cores and 64 threads, 256 GB RAM, and 8x 2.4 TB SAS 2.5-inch 6 Gbps hard drives arranged in a RAID 6 configuration. The PI will take the lead in configuring, deploying, and maintaining the server as part of his regular duties while mentoring two graduate students (Alex Moxon and Michael Vossen) in administering the server and software.

Unidata funding is requested to support this server in the absence of local funding. The university has incurred substantial unbudgeted costs and revenue losses due to the COVID-19 pandemic that furloughs and funding from the CARES and ARP Acts are insufficient to address. The College of Letters and Science suspended its technology upgrade program seven years ago due to reductions in state aid and tuition funding. Finally, discretionary funding in the Department of Mathematical Sciences (within which the Atmospheric Science Program is embedded) is severely limited (under $25,000 this year for all department activities). Consequently, no internal funding is available now, nor is expected to be available in the near future, to support this request.

b. Contributions to Unidata Community Capabilities and Broadening the Unidata Community Scope and Capabilities

The requested equipment will enable UWM to maintain its Unidata-related capabilities that are in danger of soon being lost in the absence of this funding. Through the hardware purchased in 2012, we actively participate in Unidata's IDD program and share research and operational data with the community through our THREDDS and RAMADDA servers; however, as noted above, this server has experienced multiple hard drive failures and will likely experience further component failures as it continues to age. The proposed equipment, intended to replace this server, will allow for our continued participation in the IDD and continued maintenance of our THREDDS server, allowing us to share data effectively and reliably (including new datasets, such as the four-times-daily local WRF-ARW simulations conducted by Innovative Weather) with the community for years to come.

The requested equipment will also enable UWM to deploy a local JupyterHub instance in support of educational activities. Similar to the AWIPS-II Introduction Guide that we prepared and shared with the Unidata community in 2016, we plan to document and share with the Unidata community the workflow associated with our JupyterHub deployment to provide other programs with a clear roadmap for establishing a similar local resource. In so doing, we intend to build on the successful local JupyterHub deployments at Oregon State University and Central Michigan University funded by Unidata in 2016 and 2017, respectively.

c. Contributions to Education

As an urban access institution, many of our students come from working-class backgrounds. Many of them work full-time jobs out of necessity to support themselves through college. Consequently, not all of our students can afford powerful home computers with fast, reliable Internet connections. Even for those who can, their computers are powered by a wide range of processors (e.g., Intel and ARM) and operating systems (e.g., Windows and Mac OS). Students who work full-time jobs do not have regular access to campus and our computer lab outside of regular hours. Furthermore, our students' technological sophistication varies widely, and historically our faculty have had to spend substantial time in class and during office hours working individually with students to help install Python, MATLAB, JMP Pro, and associated packages so that they can complete coursework that utilizes these programs. These factors limit the extent to which we have been able to incorporate data-driven activities within our curricula, particularly at the undergraduate level.

We intend to use the requested equipment to accelerate the deployment of Python-based scientific workflows throughout our curriculum, with an initial focus on four courses for which PI Evans has teaching responsibility, and for Innovative Weather. This will be facilitated by a local JupyterHub instance powered by The Littlest JupyterHub running on Ubuntu. JupyterHub instances facilitate data-proximate visualization and analysis, negating the need for students to each download large datasets on sometimes unreliable broadband connections to complete course activities. Students can access a JupyterHub instance from nearly any web browser, eliminating the need for them to have access to a powerful computer at home or on campus to complete course activities. Further, the server-side software installation means that students do not need to be proficient in installing Python distributions or packages and can instead focus on using Python to generate meteorological understanding. We plan to document and share with the broader Unidata community the workflow associated with our local JupyterHub deployment to provide other programs with a clear roadmap to establishing a similar resource of their own.

I. Inside the Classroom

As noted earlier in this proposal, our program is replacing the computer science course requirement in the undergraduate curriculum with a Python course effective in fall 2021. This change will allow us to integrate Python throughout our curriculum without sacrificing courses' disciplinary learning objectives or requiring that students complete more than the 120 credits required to graduate. We have identified an initial set of four courses in our undergraduate and graduate programs into which to implement a local JupyterHub instance: Synoptic Meteorology I/II, Mesoscale Meteorology, and Numerical Weather Prediction. Specific details of how the JupyterHub instance will facilitate student learning in each course are provided below:

- **Synoptic Meteorology I/II**: These core undergraduate courses, offered every other year, apply physical and dynamical principles learned in other classes to generate understanding of synoptic-scale, midlatitude meteorological phenomena. Each course is accompanied by ten to twelve lab assignments, which deepen students' understanding by through applying fundamental concepts developed during class to real data analysis. The deployment of a local JupyterHub instance will enable us to instead require that the students generate maps of their own (e.g., using real-time data that we obtain via the IDD and that they access via THREDDS and siphon) to interpret and analyze. Unidata's Python Training resource will be used to facilitate the initial development of Python notebooks that students can edit and run to prepare the required maps.
- **Mesoscale Meteorology**: This undergraduate/graduate elective, offered every other year, engages students with the physics and dynamics of midlatitude mesoscale meteorological phenomena such as lake-effect snow, orographic precipitation, gravity waves, mesoscale convective systems, and supercell thunderstorms. This course is taught as a hybrid lecture-lab course, with labs typically designed as case-study analyses of historical meteorological events using archived observations, reanalyses, and model forecasts. The deployment of a local JupyterHub instance will allow us to require that students generate maps and analyses from these data, as in Synoptic Meteorology I/II, without requiring students to individually download the necessary data.
- **Numerical Weather Prediction**: This required graduate course, offered every other year, develops students' abilities to conceptually describe and practically evaluate the strengths and weaknesses of the numerical algorithms, physical process parameterizations, and data assimilation methods used in modern numerical weather prediction models. These abilities are developed using simplified examples (Fourier analysis, 1-D advection and advection-diffusion models, 1-D ensemble data assimilation) and full-physics numerical simulations

conducted using the Advanced Research Weather Research and Forecasting (WRF-ARW) model. The local JupyterHub instance will be primarily used to support the simpler model activities, which occur both in class and in homework assignments.

II. Outside the Classroom

UWM's Innovative Weather program trains students, who are selected for the program through a competitive application process, in the rigors of operational meteorology to provide impact-based forecasts and decision support to public- and private-sector clients across the midwestern United States. Over the last nine years, we have used our Unidata-funded Dell PowerEdge T710 server to train students in data visualization to support Innovative Weather operations using GEMPAK, with selected examples available on Innovative Weather's website. The proposed JupyterHub instance will be used as a development platform to accelerate Innovative Weather student interns' curricular training in using Python tools – including Unidata's siphon and MetPy packages for data analysis – for meteorological data visualization. This will allow students to develop a robust visualization portfolio, accessible via public Jupyter Notebooks, to present to prospective employers, potentially accelerating their transition into the modern meteorological workforce.

d. Contributions to Research

Multiple UWM faculty have used our existing THREDDS server to share research data with their collaborators and the larger research community. With federal funding agencies such as NSF and NOAA and publishers such as AMS and AGU increasingly requiring that research products (such as data and code) be retained and shared with the community, the internal demand for this resource is expected to increase. This local resource fills a valuable niche, as data volumes often exceed the limits of external repositories (e.g., 50 GB at Zenodo) even when best practices for data retention are followed. Other resources such as XSEDE's TACC Ranch and UCAR's Campaign Storage are typically limited to data produced on their associated supercomputers, are not necessarily directly accessible by external users (Ranch), and/or give lowest priority to requests for storage to support funding agency or publisher data retention requirements (Campaign Storage). The requested server will enable us to maintain if not grow these contributions to research over the next five to ten years.

D. Budget and Justification

To support the proposed activities, funding is requested to purchase a Dell PowerEdge T640 server with two Intel Xeon Gold 5218 processors, 256 GB RAM, and eight 2.4 TB hard drives in a RAID 6 configuration. This configuration was designed in consultation with the Unidata Program Center and will support UWM's continued participation in the IDD, continued sharing of IDD and locally generated data with the community using THREDDS, and training students using modern Python-based scientific workflows. It is designed to support multiple simultaneous local and remote users. A cost estimate for the requested hardware, obtained from Dell Premier in March 2021, is attached and summarized below:

Quantity	Description	Total Cost
1	**Dell PowerEdge T640 Server** 2 x Intel Xeon Gold 5218 (16C/32T 2.3 GHz) Processors	$10,672

4 x 64 GB RDIMM (3200 MT/s) RAM
8 x 2.4 TB 10,000 RPM SAS 2.5” Hard Drives
PERC H730P Adapter RAID Controller, 2 GB NV Cache

Total: $10,672

This request does not include facilities and administration costs, which are not charged for capital equipment in excess of $5,000. The PI will donate time (estimated at two weeks over the full year) to configure, deploy, and maintain the requested hardware and associated software.

E. Project Milestones

The activities proposed herein will occur on the following timeline:

- **June 2021**: Hardware requote, purchase, and delivery.
- **July 2021**: Software (LDM, THREDDS, JupyterHub, etc.) installation and configuration.
- **August 2021**: Soft deployment to a small group of trusted users for testing and evaluation.
- **September 2021**: Full LDM and THREDDS deployment.
- **September 2021**: JupyterHub deployment in Numerical Weather Prediction for Fall 2021.
- **April-May 2022**: Prepare the final report, focusing on our JupyterHub deployment and the hardware’s educational benefits in the classroom and for our Innovative Weather program.

UWM offers its core atmospheric science courses only once every two years, and the other courses we have identified to use JupyterHub (Synoptic Meteorology I/II and Mesoscale Meteorology) are next scheduled to be offered during the 2022-23 academic year. Thus, JupyterHub deployment for these courses will occur beyond the period of this award.

F. References

Roebber, P. J., M. Westendorf, and G.R. Meadows, 2010: Innovative Weather: A new strategy for student, university, and community relationships. *Bull. Amer. Meteor. Soc.*, **91**, 877-888.

PowerEdge T640 Tailor Made Promo

Ships in 5 - 7 business days

Starting Price: $25,841.00 | Savings: $15,168.67

Price: $10,672.33

PowerEdge T640 Tailor Made Promo

Unit Price:	$10,672.33
Instant Savings:	$15,168.67
Order Code:	PE_T640_TM

Option	Selection	SKU/Product Code	Quantity
PowerEdge T640	PowerEdge T640 Server	210-AMBC	1
Trusted Platform Module	No Trusted Platform Module	461-AADZ	1
Chassis Configuration	Chassis with up to 16 x 2.5" SAS/ SATA Hard Drives, Tower Configuration	321-BCXG	1
Shipping	PowerEdge T640 Shipping	340-BSGU	1
Shipping Material	PowerEdge T640 Shipping Material	340-COPO	1
Regulatory	PowerEdge T640 CE, CCC, BIS Marking	389-DSUY	1
Processor	Intel® Xeon® Gold 5218 2.3G, 16C/32T, 10.4GT/s, 22M Cache, Turbo, HT (125W) DDR4-2666	338-BRVH	1
Additional Processor	Intel® Xeon® Gold 5218 2.3G, 16C/32T, 10.4GT/s, 22M Cache, Turbo, HT (125W) DDR4-2666	338-BRVH, 379-BDCO	1
Processor Thermal Configuration	Standard HS for Less = 150W,Qty 2	412-AAJW, 412-AAJW	1
Memory DIMM Type and Speed	3200MT/s RDIMMs	370-AEVR	1
Memory Configuration Type	Performance Optimized	370-AAIP	1
Memory Capacity	64GB RDIMM, 3200MT/s, Dual Rank	370-AEVP	4
RAID Configuration	C9, RAID 6 for 4 or more HDDs or SSDs (Matching Type/Speed/Capacity)	780-BCHI	1
RAID/Internal Storage Controllers	PERC H730P Adapter RAID Controller, 2GB NV Cache	405-AANW	1
Hard Drives	2.4TB 10K RPM SAS 12Gbps 512e 2.5in Hot-plug Hard Drive	400-AUTO	8
Operating System	No Operating System	619-ABVR	1
OS Media Kits	No Media Required	421-5736	1
Embedded Systems Management	iDrac9, Enterprise	385-BBKT	1
Group Manager	iDRAC Group Manager, Disabled	379-BCQY	1
Password	iDRAC,Factory Generated Password	379-BCSF	1
Additional Network Cards	On-Board Dual-Port 10GbE LOM	542-BBCT	1
Internal Optical Drive	DVD +/-RW, SATA, Internal for x8/x18/x16 chassis	429-ABCJ	1
Fans	No Additional Mid Fan for T640	384-BBSO	1
Power Supply	Dual, Hot-plug, Redundant Power Supply (1+1), 1100W	450-ADWM	1
Power Cords	NEMA 5-15P to C13 Wall Plug, 125 Volt, 15 AMP, 10 Feet (3m), Power Cord, North America	450-AALV	2
Bezel	No Bezel	325-BCON, 350-BBBW	1
Quick Sync 2 (At-the-box mgmt)	No Quick Sync for Tower	350-BBLM	1
BIOS and Advanced System Configuration Settings	Power Saving Dell Active Power Controller	750-AABF	1
Advanced System Configurations	UEFI BIOS Boot Mode with GPT Partition	800-BBDM	1
Rack Rails	No Rack Rails, No Cable Management Arm, No Casters	770-BBCR	1
System Documentation	No Systems Documentation, No OpenManage DVD Kit	631-AACK	1
Dell Services: Hardware Support	5 Years Basic Hardware Warranty Repair: 5x10 HW-Only, 5x10 NBD Onsite	815-5143, 815-5145	1
Deployment Services	No Installation	900-9997	1

Ultrabook, Celeron, Celeron Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, vPro Inside, Xeon, Xeon Phi, Xeon Inside, and Intel Optane are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

Same day shipment subject to order size limitations, Dell standard shipping methods and payment via credit card, gift card or Dell Business Credit. Notification will be

provided if there are payment delays which could impact shipping date. Electronics and accessories may ship separately.

Smart Selection. Limited quantities. Only available for orders placed by 5:59 p.m. CT Mon.–Thurs. Systems shipped the next business day after an order is placed. Subject to order approval. Software and accessories not part of the configuration will be shipped separately and may arrive after your system. Please note that Smart Selection Configuration pricing cannot be combined with other pricing offers or discounts provided or agreed to by Dell. ** Orders with Custom Factory Integration might require additional processing time.

^Dell Business Credit: Offered to business customers by WebBank, Member FDIC, who determines qualifications for and terms of credit. Taxes, shipping and other charges are extra and vary. Minimum monthly payments are the greater of $15 or 3% of the new balance shown on the monthly billing statement. Dell and the Dell logo are trademarks of Dell Inc.

**Payment solutions provided and serviced by Dell Financial Services L.L.C. or its affiliate or designee ("DFS") for qualified customers. Offers may not be available or may vary in certain countries. Where available, offers may be changed without notice and are subject to product availability, credit approval, execution of documentation provided by and acceptable to DFS, and may be subject to minimum transaction size. Offers not available for personal, family or household use. Dell and the Dell logo are trademarks of Dell Inc. Restrictions and additional requirements may apply to transactions with governmental or public entities.

[1]Subject to applicable law and regulations.

ALL ORDERS ARE SUBJECT TO APPROVAL AND ACCEPTANCE BY DELL. Pricing, availability and other terms of offer may be withdrawn or changed without notice. Dell cannot be held responsible for errors in typography or photography.