

Unidata Community Equipment Awards Cover Sheet

Proposal Title: Iowa State University / Unidata Community Equipment Proposal

Date: March 19, 2021

Principal Investigator Name: Dr. Richard Cruse
Title: Professor, Department of Agronomy
Institution: Iowa State University
Telephone number: (515) 294-7850
FAX number: 515-294-3163
Street Address:

Department of Agronomy
3202 Agronomy Hall
716 Farm House Lane
Ames, IA 50011-1051

Email address: rmc@iastate.edu

Signature of PI: _____




Principal Investigator Name: Daryl Herzmann
Title: Systems Analyst III
Institution: Iowa State University
Telephone number: (515) 294-5978
FAX number: 515-294-3163
Street Address:

Department of Agronomy
2104 Agronomy Hall
716 Farm House Lane
Ames, IA 50011-1051

Email address: akrherz@iastate.edu

Signature of PI: _____



Name of Institution Official (usually dept. chair): Dr. Kendall Lamkey

Title: Chair, Department of Agronomy

Telephone number: (515) 294-7840

FAX number: 515-294-3163

Email address: krlamkey@iastate.edu

Signature of University Official: _____




Name of Institution Official: Nichole Richter

Title: Pre-Award Administrator, Office of Sponsored Programs Administration

Telephone number: (515) 294-5225

Email address: egrants@iastate.edu

Signature of University Official: _____



Digitally signed by
Nichole Richter
Date: 2021.03.19
10:35:53 -05'00'

1. Project Summary

The Iowa State University (ISU) Meteorology and Agricultural Meteorology programs have enjoyed a long and fruitful collaboration with Unidata and the Unidata Community. Our programs extensively utilize Unidata technologies to support our teaching, research and outreach efforts. We are thankful for previous Unidata Equipment Awards granted to ISU that helped establish OpenGIS services for the Iowa Environmental Mesonet (IEM) (2005) [1] and create the “mtarchive” Thredds/HTTP service (2010, 2012) [2]. Both of these services continue to thrive. This proposal seeks to enhance and expand various IEM web services by purchasing a Dell R7525 (AMD EPYC) database server to provide two orders of magnitude more capacity.

2. Project Description

The IEM website is a heavily trafficked service providing unique data services supporting teaching and research for ISU, Unidata’s community, and beyond. While the website’s name hints at a limited geographical domain (Iowa), the site actually contains a treasure trove of curated datasets around the globe reflecting over 20 years of active development by the project’s programmer (Daryl Herzmann). Feature additions to the website are heavily guided by feedback received by Unidata Community Members. Some of these internet-unique datasets include:

- 1929-present archive of airport weather station (ASOS) data from around the globe
- National Weather Service (NWS) Watch, Warnings, and Advisories since 2005
- Massive archive of raw text products issued by the NWS since 1983
- RAOB (sounding) archive since 1946
- Billions of hydrologic observations processed from the NWS SHEF encoded feed since 2010.

These datasets and more are stored within a PostgreSQL relational database allowing for various web applications and web services to materialize data on-demand based on the user’s needs. The IEM server farm consists of a fleet of web, application, caching, and database servers. The farm has been cobbled together over the years with various funding mechanisms, so the five current database servers range in performance. For example, the ASOS and SHEF database servers have software throttles placed on the inbound web service request load to keep from being overrun.

The server proposed here would increase the IEM’s backend database capacity by two orders of magnitude due to the increased input/output operations per second (IOPS). The present fastest IEM database server (Dell R630 Intel 6410 SSDs) is able to sustain ~200k IOPS. The proposed server as configured with NVMe drives, should be able to sustain ~2 million IOPS along with lower latencies.

This **proposal advances remote teaching and classroom interaction for the geo-sciences** by increasing the speed and capacity of the IEM web services to deliver raw and aggregated data on-demand. A common usage is for the inclusion of these services within python notebooks via libraries like Unidata’s siphon to materialize a long time series of observational data for further analysis and plotting. Since all these services are openly and freely available on the Internet, the proposal funded server would have **equitable access** to all. In addition to programmatic web-services, the IEM website offers point-and-click download interfaces for various datasets which also ensures more **equitable access** for those less comfortable with writing code and just want spreadsheets.

The IEM curates an extensive archive of hydrologic data provided from the NWS in SHEF format. Much of this data is outside the typical bailiwick of meteorology, and it is additionally very difficult to

find archives of these data. IEM implementation of web services for this dataset are limited by database capacity, and this proposal would further open this **locally-created dataset and derived products to the wider Unidata community.**

2a. Technical Implementation of the Server

The software setup of this server would closely follow some excellent documentation [3] provided by the “Let’s Encrypt” project, which recently setup MySQL on this very same server hardware. Their Dell R7525 was loaded to the gills with expensive components, but the server with this proposal should perform similarly. OpenZFS on Red Hat Enterprise Linux 8 would be used to provide fault tolerance and scalability for the NVMe drives, since hardware RAID is not an option. OpenZFS would also allow compression to be used at the filesystem level, which would greatly increase (2-3x from local testing) the amount of data it could store. **Project PI Herzmann has extensive server administration and database software configuration experience, so no additional support help should be necessary.** He is permitted to administer the equipment provided by this proposal.

The server would reside within the present IEM computing cluster, which spans two data centers on the ISU Campus providing some limited redundancy. Compute intensive jobs are run within a virtual machine Red Hat Enterprise Virtualization (RHEV) environment that is managed centrally by the college. “Bare-metal” servers, like database servers, are permitted within this environment and are communicated with over isolated VLANs and not necessarily managed by RHEV.

Proposing a single server likely raises concerns about fault tolerance and what happens if this server should die. The server has a five year warranty with four hour parts, so that provides some peace of mind. For other types of failures, the IEM utilizes PostgreSQL replication to have a hot standby ready to take on writes should the primary server fail. The write performance necessary for the IEM real-time processing is trivial, so a warm standby with significantly lesser hardware will work. A lengthy fail-over would require very tight throttles to keep the webfarm load from over-running the backup hardware, but the main goal is to not have data loss nor missed real-time ingest.

2b. Addressing proposal review criteria

1. **Intellectual or technical merit:** For many of the datasets mentioned in this proposal, a relational database server remains the best technical solution to run on limited and commodity hardware. This approach has worked very well for the IEM over the past 20 years of existence and the choice of PostgreSQL + PostGIS remains actively developed with useful new features being added with each release. The hardware requested would not preclude the implementation of some cutting-edge technology like Timescale-DB or No-SQL solution, if the need arises.
2. **Contribution to Unidata capabilities:** Unidata has been a pioneer in remote / data proximate software systems with libraries like siphon and servers like Thredds helping to provide data to users as-needed and on-demand. This proposal would enhance the sounding archive service used within siphon and the proposal PI commits to helping to implement the other archives mentioned within siphon. Truth be told, the PI would like to do that anyway regardless of proposal funding, but this would certainly provoke more Github pull requests!
3. **Broadening Unidata scope and capabilities:** The IEM archives of primarily point data fill in a gap that is not well facilitated by Thredds, which is more geared toward grid and short time series data. The new services exposed within siphon would additionally help MetPy users wanting such data, but finding it lacking within Unidata Thredds servers.

4. **Contributing to the advancement of technology:** This proposal does not really introduce any new technical approach to the classic problem of how to provide data over the web. This just builds more capacity to support more intensive web services. As one example, a user could request “provide me all CONUS growing degree day data with a provided custom base and ceiling along with the climatology between 15 April and 18 August”. The IEM presently cannot provide this as the backend database would take a longer time to compute than practical for web sessions.
5. **Contribution to education:** Python notebooks are heavily used at ISU and other Unidata-affiliated Universities. Improvements to speed and breath of web services provided will directly support these notebooks.
6. **Contribution to research:** Experience has shown researchers use whatever technology stack they are most comfortable with. The web services supported by this proposal are generally agnostic, so researchers should benefit as well by having access to real-time and massive archives of tough to find datasets.

3. Budget

The included Dell quote for the single server comes in at \$13,130.37. ISU has agreed to waive indirect costs. Server installation, setup, and administration will be covered by proposal PI with datacenter costs of housing the server covered by other funding streams.

The budgeted server was goal-seeked to come well underneath the Unidata \$20K limit while still fulfilling the goals outlined within the proposal. Hopefully this allows for more flexibility for the number of proposals that could be funded. If there is more money available after the funded proposals by Unidata are decided, we certainly could add more NVMe + memory to this proposal to fulfill even more goals! :) We’ll need to redo the quoting process anyway as the Dell quote will have expired by the time any funding is awarded. Regardless of funding availability, any re-quoted server would not exceed the amount budget here.

4. Project Time-line

We would effort the purchase and deployment of the funded server as soon as possible in order to meet the present needs outlined within the proposal. Space is currently available for our allocation within the data center, so it should be up and running within days of receipt and some performance bench marking + testing. An approximate time line is as follows:

- May 2021: Proposal awarding and accounting book work.
- Early June 2021: Purchase and delivery of server by Dell.
- Late June 2021: Server provisioned and relevant databases migrated to it.
- May-Sep 2021: Software development of additional web services this new hardware can support. Submission of pull requests to siphon to expose some current + new services there.
- September 2021: Report submitted to Unidata

[1] <https://mesonet.agron.iastate.edu/docs/unidata2006/>

[2] <https://mtarchive.geol.iastate.edu/thredds/catalog.html>

[3] <https://letsencrypt.org/2021/01/21/next-gen-database-servers.html>



A quote for your consideration.

Based on your business needs, we put the following quote together to help with your purchase decision. Below is a detailed summary of the quote we've created to help you with your purchase decision.

To proceed with this quote, you may respond to this email, order online through your [Premier page](#), or, if you do not have Premier, use this [Quote to Order](#).

| | | | |
|------------------|------------------------|-------------------|--------------------------|
| Quote No. | 3000077185505.1 | Sales Rep | Jessica Nunn |
| Total | \$13,130.37 | Phone | (800) 456-3355, 6180057 |
| Customer # | 145407962 | Email | Jessica_Nunn@DELL.com |
| Quoted On | Jan. 19, 2021 | Billing To | CORY HARMS |
| Expires by | Feb. 14, 2021 | | IOWA STATE UNIV |
| Solution ID | 13635551 | | 1520 ADMIN SERVICES BLDG |
| Deal ID | 17256076 | | IOWA STATE UNIVERSITY |
| | | | AMES, IA 50011 |

Message from your Sales Rep

Please contact your Dell sales representative if you have any questions or when you're ready to place an order. Thank you for shopping with Dell!

Regards,
Jessica Nunn

Shipping Group

| Shipping To | Shipping Method |
|--------------------------------------------------------------------------------------------------------------|-------------------|
| DARYL HERZMANN IOWA STATE UNIV 2104 AGRONOMY HALL AGRONOMY AMES, IA 50011-1012 (515) 294-4860 | Standard Delivery |

| Product | Unit Price | Qty | Subtotal |
|-----------------------------------|-------------|-----|-------------|
| PowerEdge R7525 Database Server | \$13,130.37 | 1 | \$13,130.37 |

| | |
|----------------------------|--------------------|
| Subtotal: | \$13,130.37 |
| Shipping: | \$0.00 |
| Non-Taxable Amount: | \$13,130.37 |
| Taxable Amount: | \$0.00 |
| Estimated Tax: | \$0.00 |

| | |
|---------------|--------------------|
| Total: | \$13,130.37 |
|---------------|--------------------|

Consult your DFS rep. Special lease pricing may be available.

Shipping Group Details

Shipping To

DARYL HERZMANN
IOWA STATE UNIV
2104 AGRONOMY HALL
AGRONOMY
AMES, IA 50011-1012
(515) 294-4860

Shipping Method

Standard Delivery

| | | Qty | Subtotal |
|------------------------------------------|--------------------|----------|--------------------|
| PowerEdge R7525 Database Server | \$13,130.37 | 1 | \$13,130.37 |

Estimated delivery if purchased today:
Mar. 02, 2021
Contract # C000000181059

| Description | SKU | Unit Price | Qty | Subtotal |
|----------------------------------------------------------------------------------|----------|------------|-----|----------|
| SAS/SATA/NVMe Capable Backplane | 379-BDSW | - | 1 | - |
| PowerEdge R7525 Server | 210-AUVQ | - | 1 | - |
| No Trusted Platform Module | 461-AADZ | - | 1 | - |
| 16X 2.5 SAS/SATA + 8X 2.5 NVME without XGMI | 321-BFEE | - | 1 | - |
| AMD 7302 3GHz,16C/32T,128M,155W,3200 | 338-BSWC | - | 1 | - |
| AMD 7302 3GHz,16C/32T,128M,155W,3200 | 338-BSWC | - | 1 | - |
| Additional Processor Selected | 379-BDCO | - | 1 | - |
| Heatsink for CPU less than 180W | 412-AASF | - | 1 | - |
| Heatsink for CPU less than 180W | 412-AASF | - | 1 | - |
| Plastic Shroud | 750-ACFV | - | 1 | - |
| Performance Optimized | 370-AAIP | - | 1 | - |
| 3200MT/s RDIMMs | 370-AEVR | - | 1 | - |
| Unconfigured RAID | 780-BCDS | - | 1 | - |
| PERC H745 Controller, Front | 405-AAUZ | - | 1 | - |
| Front PERC Mechanical Parts, for 2.5" x8, x10, x16 SAS/SATA or 2.5" NVMe Chassis | 750-ACFQ | - | 1 | - |
| Performance BIOS Settings | 384-BBBL | - | 1 | - |
| UEFI BIOS Boot Mode with GPT Partition | 800-BBDM | - | 1 | - |
| Standard Fan x6 | 750-ACSP | - | 1 | - |
| Dual, Hot-Plug,Power Supply Redundant (1+1), 1400W ART | 450-AJHI | - | 1 | - |
| Riser Config 3, Half Length, 5 x16 slots | 330-BBPI | - | 1 | - |
| Assembly BOSS Blank R7525 | 329-BERC | - | 1 | - |
| R7525 Motherboard with LOM | 384-BCKH | - | 1 | - |
| Software System R7525 15G | 384-BCKI | - | 1 | - |
| iDRAC9,Enterprise x5 | 385-BBOT | - | 1 | - |
| Broadcom 5720 Quad Port 1GbE BASE-T Adapter, OCP NIC 3.0 | 540-BCOB | - | 1 | - |
| PowerEdge 2U Standard Bezel | 325-BCHU | - | 1 | - |
| No Quick Sync | 350-BBYX | - | 1 | - |
| iDRAC,Legacy Password | 379-BCSG | - | 1 | - |
| iDRAC Group Manager, Enabled | 379-BCQV | - | 1 | - |

| | | | | |
|-----------------------------------------------------------------------------------------------------------------------------------------------|----------|---|----|---|
| No Operating System | 611-BBBF | - | 1 | - |
| No Media Required | 605-BBFN | - | 1 | - |
| OpenManage Integration for VMware vCenter - 1 host increment, 5 year license - Digitally Fulfilled | 634-BJBC | - | 1 | - |
| ReadyRails Sliding Rails | 770-BBBQ | - | 1 | - |
| No Systems Documentation, No OpenManage DVD Kit | 631-AACK | - | 1 | - |
| PowerEdge R7525 Shipping | 340-COFR | - | 1 | - |
| R7525 Shipping Material | 481-BBFE | - | 1 | - |
| PowerEdge R7525 CE Marking , No CCC Marking | 389-DUZE | - | 1 | - |
| ProSupport Plus Mission Critical 7x24 HW-SW Technical Support and Assistance 5 Years | 829-7633 | - | 1 | - |
| ProSupport Plus Mission Critical 4-Hour 7x24 On-Site Service with Emergency Dispatch 3 Years | 829-7644 | - | 1 | - |
| ProSupport Plus Mission Critical 4-Hour 7x24 On-Site Service with Emergency Dispatch 2 Years Extended | 829-7646 | - | 1 | - |
| Dell Hardware Limited Warranty Plus On-Site Service | 829-7738 | - | 1 | - |
| Thank you for choosing Dell ProSupport Plus. For tech support, visit //www.dell.com/contactdell | 951-2015 | - | 1 | - |
| Dell Limited Hardware Warranty Plus Service, Extended Year(s) | 975-3462 | - | 1 | - |
| On-Site Installation Declined | 900-9997 | - | 1 | - |
| 8GB RDIMM, 3200MT/s, Single Rank | 370-AEVO | - | 16 | - |
| 480GB SSD SATA Mix Use 6Gbps 512 2.5in Hot-plug AG Drive, 3 DWPD, 2628 TBW | 400-AZUT | - | 2 | - |
| 1.6TB Enterprise NVMe Mixed Use AG Drive U.2 Gen4 with carrier | 400-BKGF | - | 5 | - |
| Jumper Cord - C13/C14, 0.6M, 250V, 13A (North American, Guam, North Marianas, Philippines, Samoa) | 492-BBDH | - | 2 | - |

| | |
|-----------------------|--------------------|
| Subtotal: | \$13,130.37 |
| Shipping: | \$0.00 |
| Estimated Tax: | \$0.00 |
| Total: | \$13,130.37 |

Important Notes

Terms of Sale

This Quote will, if Customer issues a purchase order for the quoted items that is accepted by Supplier, constitute a contract between the entity issuing this Quote ("Supplier") and the entity to whom this Quote was issued ("Customer"). Unless otherwise stated herein, pricing is valid for thirty days from the date of this Quote. All product, pricing and other information is based on the latest information available and is subject to change. Supplier reserves the right to cancel this Quote and Customer purchase orders arising from pricing errors. Taxes and/or freight charges listed on this Quote are only estimates. The final amounts shall be stated on the relevant invoice. Additional freight charges will be applied if Customer requests expedited shipping. Please indicate any tax exemption status on your purchase order and send your tax exemption certificate to Tax_Department@dell.com or ARSalesTax@emc.com, as applicable.

Governing Terms: This Quote is subject to: (a) a separate written agreement between Customer or Customer's affiliate and Supplier or a Supplier's affiliate to the extent that it expressly applies to the products and/or services in this Quote or, to the extent there is no such agreement, to the applicable set of Dell's Terms of Sale (available at www.dell.com/terms or www.dell.com/oemterms), or for cloud/as-a-Service offerings, the applicable cloud terms of service (identified on the Offer Specific Terms referenced below); and (b) the terms referenced herein (collectively, the "Governing Terms"). Different Governing Terms may apply to different products and services on this Quote. The Governing Terms apply to the exclusion of all terms and conditions incorporated in or referred to in any documentation submitted by Customer to Supplier.

Supplier Software Licenses and Services Descriptions: Customer's use of any Supplier software is subject to the license terms accompanying the software, or in the absence of accompanying terms, the applicable terms posted on www.Dell.com/eula. Descriptions and terms for Supplier-branded standard services are stated at www.dell.com/servicecontracts/global or for certain infrastructure products at www.dellemc.com/en-us/customer-services/product-warranty-and-service-descriptions.htm

Offer-Specific, Third Party and Program Specific Terms: Customer's use of third-party software is subject to the license terms that accompany the software. Certain Supplier-branded and third-party products and services listed on this Quote are subject to additional, specific terms stated on www.dell.com/offeringspecificterms ("Offer Specific Terms").

In case of Resale only: Should Customer procure any products or services for resale, whether on standalone basis or as part of a solution, Customer shall include the applicable software license terms, services terms, and/or offer-specific terms in a written agreement with the end-user and provide written evidence of doing so upon receipt of request from Supplier.

In case of Financing only: If Customer intends to enter into a financing arrangement ("Financing Agreement") for the products and/or services on this Quote with Dell Financial Services LLC or other funding source pre-approved by Supplier ("FS"), Customer may issue its purchase order to Supplier or to FS. If issued to FS, Supplier will fulfill and invoice FS upon confirmation that: (a) FS intends to enter into a Financing Agreement with Customer for this order; and (b) FS agrees to procure these items from Supplier. Notwithstanding the Financing Agreement, Customer's use (and Customer's resale of and the end-user's use) of these items in the order is subject to the applicable governing agreement between Customer and Supplier, except that title shall transfer from Supplier to FS instead of to Customer. If FS notifies Supplier after shipment that Customer is no longer pursuing a Financing Agreement for these items, or if Customer fails to enter into such Financing Agreement within 120 days after shipment by Supplier, Customer shall promptly pay the Supplier invoice amounts directly to Supplier.

Customer represents that this transaction does not involve: (a) use of U.S. Government funds; (b) use by or resale to the U.S. Government; or (c) maintenance and support of the product(s) listed in this document within classified spaces. Customer further represents that this transaction does not require Supplier's compliance with any statute, regulation or information technology standard applicable to a U.S. Government procurement.

For certain products shipped to end users in California, a State Environmental Fee will be applied to Customer's invoice. Supplier encourages customers to dispose of electronic equipment properly.

Electronically linked terms and descriptions are available in hard copy upon request.

^Dell Business Credit (DBC):

OFFER VARIES BY CREDITWORTHINESS AS DETERMINED BY LENDER. Offered by WebBank to Small and Medium Business customers with approved credit. Taxes, shipping and other charges are extra and vary. Minimum monthly payments are the greater of \$15 or 3% of account balance. Dell Business Credit is not offered to government or public entities, or business entities located and organized outside of the United States.