

Agenda: Spring 2024 Users Committee Meeting

NSF Unidata, UCAR Foothills Lab 1, EOL Atrium

(Times are Mountain Daylight Time)

Monday, 13 May 2024

08:30 – 09:00 Doors Open

09:00 – 9:30 Welcome and Administrative Items (Casey Davenport & Alex Davies / Tanya Vance)

- Date for Fall Joint Committee meeting
- Updates from Members (Committee)
- Updates from UsersCom Chairs from Spring SAC Meeting

09:30 – 10:30 Staff Status Reports (All)

- Visualization Software/Tools
- Data Access/Formats/Dissemination
- Community Services/Educational Efforts

10:30 – 10:45 Break

10:45 – 11:00 UPC Share Out

11:00 – 12:00 Director's Report (Mohan Ramamurthy)

12:00 – 13:30 Lunch at the Foothills Lab Cafeteria

13:30 – 14:00 NCEP Report and Questions (Tony Salemi)

14:00 – 15:30 Interactive Committee Brainstorming Session

- Primer
- Break
- Increasing Committee Engagement
- Broadening Community Participation
- Reconvene for Share Out and Next Steps

15:30 – 16:00 Wrap-up Day One

16:00 Adjourn

18:30 Dinner at [Avanti](#) F&B Boulder A Collective Eatery, 1401 Pearl Street, Boulder, CO 80302 ([map](#)).

The restaurant will provide validated parking for the privately-owned parking lot at the [1434 Spruce St Garage](#). (Note that the entrance to this garage is on 15th street, which is one way going north.)

Tuesday, 14 May 2024

08:30 – 09:00 Doors Open

09:00 - 09:15 Convene and Outstanding Items from Previous Day

09:15 – 09:45 DeSouza Award. Discuss Candidates for the 2024 Honor (Committee)

09:45 – 10:30 Equipment Awards Expansion and Reporting

10:30 – 10:45 Break

10:45 – 11:00 Establishing an Ad-hoc Committee on NSF Unidata-AMS Ed Guidelines (Casey Davenport)

11:00 – 11:45 NSF Unidata Committee Member Project Shareout (Alex Davies)

11:45 – 12:30 All Other Business

- Any Items from Previous Day
- Action Items, Fall Meeting Dates

12:30 Close Meeting

Status Report: Users Committee Actions

November 2023- April 2024
Unidata Program Center Staff

Actions from the Previous Meeting (November 2023)

Action 1

Alex Davies will discuss options to add oceanographic data to data feeds with Mike Zuranski.

Result

This has not yet happened.

Action 2

Alex Davies will discuss WRF applications for the Naval Academy and Ocean Community with Jeff Weber.

Result

Channel of communication open between Alex and Jeff. Alex and Jeff will present at the upcoming Spring meeting.

Action 3

Todd Murphy and Aaron Kennedy will share information and syllabus for basic computer course information with the committee.

Result

This has not yet happened.

Action 4

Adding an agenda item for the Spring meeting to discuss Learner Profile and Learning situations when designing learning services (Tanya Vance, Nicole Corbin, Casey Davenport).

Result

This agenda item will be part of the 2024 Fall Users Committee meeting to better align with availability and the launch of the beta release of the NSF Unidata Educational Hub.

Action 5

The committee will engage their communities to seed and capture thoughts on the community hubs concept and increasing MSI partnerships.

Result

This has not yet happened.

Action 6

Mohan Ramamurthy will discuss opportunities to expand the resource allocation types for the Equipment Awards with NSF such as funding workshops and Train the Trainers, building cloud resources, etc. Tanya Vance will follow up with the committee.

Result

The discussion and opportunities to expand allocation types of the Equipment Awards is an item on the Spring 2024 agenda.

Status Reports Executive Summary

November 2023- April 2024
Unidata Program Center Staff

This summary is compiled from the full status reports, available online:
[Staff Status Reports: November 2023 - April 2024](#)

Visualization and Analysis Software and Tools

AI/ML

Unidata ML Staff have been working on educational materials, Jupyter Notebooks, and blog posts. The Cybertraining Award from NSF is a focus moving forward with our colleagues at MSU Denver for in classroom AI/ML content and materials. **This no-code module is ready to be shared at other institutions**, please get in touch if you are interested. The module is designed to be dropped into an upper level earth systems science course. Short course / workshop curriculum is also in development, to be delivered both in person and virtually.

AWIPS

Since the last status report, the team has successfully completed beta development of AWIPS version 20.3.2 and released a production build in December, 2023. Up until now, the team has also supported a v18 EDEX in case of emergencies for academic institutes that were dependent on the previous version. Immediately following the production release the team began development on version 23.4.1 of AWIPS which will migrate us from RHEL7 (and CentOS7) dependencies to RHEL 8. This is especially pressing because CentOS7 has an EOL for June 30th. The AWIPS team is aiming to put out a beta release of v23.4.1 by the time you are reading this report.

Aside from software development, the team had a presence at AMS 2024, with Tiffany Meyer hosting an AWIPS Student Workshop and a talk highlighting the updates of NSF Unidata's developments. Thank you to our friends at Texas A&M and Dr. Nowotarski for putting us in touch with meteorology student Victoria Elliott who also presented at our AMS student workshop!

Finally, the AWIPS team has been involved in some work to try and develop funding opportunities from the private and commercial sectors that are interested in using our software and services. The team hopes to make a bit of progress on this venture to help ease our upcoming budget concerns.

IDV with RAMADDA

We continue to support, update, and enhance the 3D data visualization and analysis tool IDV for our community. Our current activities include: coordinating with netCDF-Java group to add new data formats, collaborating with the SSEC developers to enhance the VisAD library, and working with our community to promote the usage of the IDV in research and education.

Python

Unidata's Python efforts continue to encompass: training on the use of Python for the community; development and maintenance of several tools for the community (most notably MetPy but also Siphon and data processing scripts); and participation within the broader scientific Python community. We continue to lead and support a variety of educational efforts, including our first collaboration with the USGS on a training session. We are also furthering development of asynchronous training materials through Project Pythia, where we are working to migrate our existing workshop and gallery materials into a so-called "cookbook" within the broader project. MetPy development continues with the 1.6.0 feature release (largely calculation additions and standardization of relative humidity definition) and 1.6.1/1.6.2 bugfix releases. We are in the early planning stage for a 1.7 release which will center around min/max identification and S3 cloud data clients. Community use of the library is proving extensive, with 90 theses and peer-reviewed publications mentioning or citing MetPy in 2023; this brings the total count to over 300. We continue to assist the broader community with participation within matplotlib, cartopy, and conda-forge, though it has become increasingly difficult to dedicate time to these efforts given the full portfolio of responsibilities.

Questions for the Committee

1. What are the biggest challenges to incorporating AI/ML into curriculum and research?
2. What python packages are the biggest source of frustration in your teaching?
3. Please let us know if you use AWIPS in the classroom, if not, what's the biggest obstacle keeping you from doing so?
4. We have noticed that many advanced features of the IDV, such as formulas and trajectory displays, have not been widely used in the community and many data servers that the IDV can directly access are less well known to IDV users. We would like to provide help to classes, research groups and project teams to use these resources. Can committee members help to establish such connections?

Data Access/Formats/Dissemination

Community Data Standards

Unidata's netCDF teams continues to engage with the Zarr community on:

- 1) Zarr support in both the netCDF-C and netCDF-Java libraries;
- 2) the development of the Zarr version 3 specification; and
- 3) the development of the GeoZarr convention.

Unidata continues to be active in efforts to advance the Climate and Forecast (CF) Conventions for netCDF.

Unidata continues to be active in several international standards bodies and other communities focused on data and technology including the World Meteorological Organization (WMO), the Open Geospatial Consortium (OGC), and the Earth System Information Partners (ESIP).

Data Services

Progress on the ~~RTSTATS Revamp~~ RTSTATS-NG has resumed. Deployment of an early release is underway and additional features are being developed.

Revising how we collect data metrics: Improving consistency and confidence in our reporting. A focus will be on our IDD and ADDE statistics collection process as a result of a recent analysis highlighting challenges & opportunities for improvement.

User support of course continues to be a primary focus. See the **Status Report: Support** for inquiry metrics.

GOES

NSF Unidata continues to operate satellite downlink facilities for the NOAA Port Satellite Broadcast Network (SBN) and GOES-East and GOES-West rebroadcast services on behalf of UCAR/NCAR and the NSF Unidata community. All received products are then provided via the Internet Data Distribution system (IDD) in various feeds and via remote access provided by AWIPS EDEX, McIDAS ADDE and THREDDS Data Servers.

IDD

Unidata continues to support, update, and enhance the data available via the IDD for the benefit of research and education. Included but not limited to adding new data formats, bridging the knowledge gap in newly introduced data, and providing statistics of data flow and composition.

IT

Our role is to maintain and enhance the productivity of the staff and assist with the resolution of issues in service to the community. Primarily, that consists of keeping end-user and developer systems secure, and keeping servers and services highly available, patched, and operational for the community. This report is informational and there are no pressing issues.

LDM

Unidata's LDM team continues to update source code and operating paradigms with ever-changing user demographics and user requirements, particularly in the area of security and inclusiveness of data.

NetCDF

The netCDF team continues to work towards maintaining the reliability of the netCDF libraries, while keeping one eye forward as to the future needs of our community. We have continued our community engagement efforts and collaborations whenever and wherever possible; examples of this include our involvement with the Zarr Community meetings and our membership on the Zarr Enhancement Protocol (ZEP) committee. We have also continued conversation with the HDF group and other community groups working on similar efforts.

We continue to address the issues associated with the proliferation of new mainstream architectures (Apple Developed M1/2/3/ARM), evolving compilers and standards, and

extending our collaborations with tangential, but related, projects (conda-forge libnetcdf feedstock, for example).

In the past several months, the NetCDF team has participated in a developer exchange program with the Atmospheric Chemistry Observations & Modeling (ACOM) at NCAR. The ACOM developer, Kyle Shores, was able to significantly modernize the build infrastructure for netCDF-C, freeing up resources for the core development team to work on ncZarr and S3 support, as well as user support and general quality-of-life technical improvements.

THREDDS

The deprecation of the Spring 5.x library sets a hard deadline for a significant number of upgrades in the TDS by December 2024. We plan to, at that time, release a long-term maintenance version of the TDS, before shifting focus to a replacement data service system.

Questions for the Committee

None at this time.

Community Services and Educational Efforts

Science Gateway & Cloud Computing Activities

- LROSE collaboration with CSU and NCAR
- WRF on Jetstream2 for classrooms and TCUs
- Development phases of Science Gateway Re-Imagined outlined
- Enhanced AWIPS EDEX servers on Jetstream2 Cloud
- JupyterHubs for atmospheric science education including GPU and Dask Hubs
- Maintained NEXRAD THREDDS server
- Updated and streamlined Docker image deployments for LDM, RAMADDA, THREDDS
- Strengthened cloud security
- Presented at Science Gateways Conference, 2023

Community Services

In addition to “normal, day-to-day” activities of communication and coordination with community members, the Community Services group’s efforts in the past six months have been focused on:

- Outreach and engagement with historically marginalized and underrecognized communities, including activities with tribal colleges and universities and the Sovereign Data Network project partners.
 - Some notable activities include live data streaming and data access from SIPI and NTU, collaborating with Lamont-Doherty Earth Observatory (LDEO) on the Hydrogeologic Impacts of Terrestrial Expressions of Mesozoic Perturbations (HI-TEMP) drilling project on the Navajo Nation, and working with The Haskell Foundation, Nebraska Indian Community College, Oglala Community College, Aaniiih Nakoda College, and Rising Voices Changing Coasts Hub to expand the Sovereign Data Network
- Outreach to the Earth System Science community through participation at the following:

- Conferences: AMS, AGU, AIHEC, Rising Voices, and Environmental Data Science Innovation & Inclusion Lab (ESIIL) Innovation Summit
- Workshops and Events: MSU Career Day, SciEd Career Highlights, the NSF NCAR Earth System Science and Technology hubs (NESTs) Community Workshop on Conceptualization and Coastal Resiliency, A New Lens for Community Engagement Workshop, and the Exploring Data Sovereignty and Sovereign Data Network Workshop
- Working Groups: NSF NCAR Convergence Science Community Network, NSF NCAR-MSI Collaborators Community of Practice, NSF SOARS, NSF Earth Data Relations, and NSF ESIL
- Expansion of learning materials, resources, and offerings
- Progressing the Reimagined Science Gateway and integrated Education Hub to make our educational services more discoverable on the web as well as promoting the variety of educational services we are able to provide, including updating our [website](#) and community [requests](#)
- Development and delivery of the first module on Machine Learning Foundations in the Earth Systems Sciences at Metropolitan State University of Denver for [NSF Award #2319979](#)
- Facilitating a MetPy AMS Short Course, Python and AWIPS Student Conference sessions, and an AMS oral presentation on data readiness microlearning
- Partnering with COMET and USGS on the development of a series of NetCDF/CF microlearning resources and delivery of the Metpy for Quantitative Analysis training Spring - Summer 2024
- Providing coordination for the 2024 NSF Unidata Community Equipment Awards
- Supporting submission and post-submission actions of NSF Unidata's next core funding award proposal
- Leading the planning and development of UCAR Community Programs (UCP)'s Strategic Plan and supporting the review of UCAR's Strategic Plan

Support

Unidata staff have started vetting replacement packages for our current eSupport system. We hope to have a new support package in use by the end of Spring 2024.

The AWIPS and THREDDS developers have performed a staggering amount of support for their packages this past year (820 and 549 support replies respectively). 🙌

Questions for the Committee

1. What new or existing software are you excited to use or teach in your research? Does the installation process pose a barrier to entry, and can the Unidata Science Gateway team help you?
2. Are you interested in being a beta tester of the Re-Imagined Science Gateway?
3. Are there any specific case studies or success stories we should highlight?
4. What feedback do you have on the phased SGRI development approach?
5. What changes/recommendations would you like to see in Unidata support or how Unidata conducts support? Are you subscribed to any of the unidata software package mailing lists?

Prepared *April 2024*

Status Report: AI/ML

November 2023- April 2024

Thomas Martin, Nicole Corbin

Executive Summary

Unidata ML Staff have been working on educational materials, Jupyter Notebooks, and blog posts. The Cybertraining Award from NSF is a focus moving forward with our colleagues at MSU Denver for in classroom AI/ML content and materials. **This no-code module is ready to be shared at other institutions**, please get in touch if you are interested. The module is designed to be dropped into an upper level earth systems science course. Short course / workshop curriculum is also in development, to be delivered both in person and virtually.

Questions for Immediate Committee Feedback

- What are the biggest challenges to incorporating AI/ML into curriculum and research?
- What python packages are the biggest source of frustration in your teaching?

Activities Since the Last Status Report

- Working hard NSF Cybertraining award! The first module is being delivered this month (April) at MSU Denver. This is now in a shareable state!
- NOAA now delivering ML model (GraphCast)
output:<https://registry.opendata.aws/noaa-aws-graphcastgfs-pds/>
- Visited Texas Tech to assess ML & python needs and priorities, along with giving a seminar talk (would love to give more!)
- Taught two day intro to ML course at NCAR, resources can be found here:
https://github.com/NCAR/ML_workshop2023/tree/main/tutorials
- Office hours still available (sign up here:
<https://calendar.app.google/ZsM8dLHLA65eGAr39>)
- Numerous blog posts, feel free to request

New Activities

We are looking for more ways to support our community at large. This might include in-person or remote workshops, 1 on 1 mentoring and tutoring, and online asynchronous resources.

The Cybertraining award is a major focus for the duration of the award (2 years). Always looking for more collaboration opportunities around ML/AI education in the classroom, from guest lectures to content development.

Prepared *April 2024*

Status Report: AWIPS

November 2023- April 2024

Tiffany Meyer, Shay Carter

Executive Summary

Since the last status report, the team has successfully completed beta development of AWIPS version 20.3.2 and released a production build in December, 2023. Up until now, the team has also supported a v18 EDEX in case of emergencies for academic institutes that were dependent on the previous version. Immediately following the production release the team began development on version 23.4.1 of AWIPS which will migrate us from RHEL7 (and CentOS7) dependencies to RHEL 8. This is especially pressing because CentOS7 has an EOL for June 30th. The AWIPS team is aiming to put out a beta release of v23.4.1 by the time you are reading this report.

Aside from software development, the team had a presence at AMS 2024, with Tiffany Meyer hosting an AWIPS Student Workshop and a talk highlighting the updates of NSF Unidata's developments. Thank you to our friends at Texas A&M and Dr. Nowotarski for putting us in touch with meteorology student Victoria Elliott who also presented at our AMS student workshop!

Finally, the AWIPS team has been involved in some work to try and develop funding opportunities from the private and commercial sectors that are interested in using our software and services. The team hopes to make a bit of progress on this venture to help ease our upcoming budget concerns.

Questions for Immediate Committee Feedback

Please let us know if you use AWIPS in the classroom, if not, what's the biggest obstacle keeping you from doing so?

Activities Since the Last Status Report

AWIPS

Our EDEX servers have been continuously running on the new Jetstream2 platform since it was released, however there have been numerous issues (from Jetstream or their vendors) that we've encountered on the new platform. The AWIPS team has worked closely with Unidata IT, Science Gateway coworkers (Ana and Julien) and the Jetstream staff to troubleshoot and resolve issues that arise. We have been maintaining a suite of EDEX servers since the beginning of the year: a production v18 server, both a production and

development/backup v20 server, and as of recently, a v23 development server. Our production server continues to serve real-time weather and geographic data to [CAVE clients](#) and the [python-awips](#) data access framework API.

Through the use of ancillary EDEX machines we have been able to decouple certain datasets from the main EDEX instance. We take advantage of distributing EDEX workload over three machines: a main EDEX, an ancillary radar EDEX, and an ancillary satellite EDEX. These [distributed architectural concepts](#) of AWIPS allow us to scale EDEX in the cloud to account for the size of incoming data feeds.

Since the last status report, we've put out two production releases of v20 and are working hard on an upcoming v23 beta release which will include only a CAVE virtual machine and linux distribution . Our latest release is [version 20.3.2-2](#). A link to all of our AWIPS release notes can be found [here](#). We have also put out a [production release of v20 python-awips](#). The production release of v20 also came with a new offering of Alaska data – we now serve NAM data for the AK region, thanks to a user request that we were able to accommodate.

Throughout the development cycles of the beta and production releases, we have been collecting user feedback through [this reporting form](#). This method has been an invaluable tool to collect, track, and update users on issues found in version 20 of AWIPS. The team plans to create a similar reporting form to use in development with the new v23 AWIPS as well.

Version 23.4.1 will have some significant upgrades in the software dependencies, the most notable being the Operating System and Python:

- RHEL (from RHEL7 to RHEL8)
- Python (from 3.6 to 3.11)
- Java (from OpenJDK 11.0.13 to 11.0.22)
- Eclipse (from 4.17.0 to 4.21.0)

The Operating System is significant because currently the NSF Unidata version of CAVE is dependent on RHEL7 (RedHat Enterprise Linux version 7) and we use CentOS7 as a compatible, free option for the operating system. Unfortunately, CentOS7 is reaching its End of Life (EOL) on June 30th, 2024, so we are pushing to have a production release of AWIPS that depends on a newer operating system. We have written an [AWIPS Tips blog](#) discussing our plans around this issue, but the summary is version 23.4.1 of AWIPS is RHEL8 dependent and we will be using Rocky8 as our free operating system. It is important to note that previous versions of AWIPS are not compatible with the new version 23.

A significant portion of our documentation both for [CAVE, EDEX](#), and [python-awips](#) has been modified for easier understanding and comprehension. For the production release of v20, all of our documentation was reviewed and updated with all new screenshots and instructions. We are continuing to update our python-awips example notebooks to follow our new


template that contains a helpful table of contents, with consistent subsections across the various example topics.

Our blog series, [AWIPS Tips](#), has successfully been running every other week for over three years now. A current list and breakdown of all the entries is provided on our [documentation website in the Educational Resources page](#). We plan to continue the blog series for the foreseeable future and have several more ideas already planned out for upcoming entries. Announcements for new blog posts are shared through our mailing list (awips2-users@unidata.ucar.edu), and our social media accounts (Facebook, Twitter, LinkedIn, and YouTube when applicable).

We have asynchronous training available for both CAVE and Python-AWIPS on the Educational Resources website. We encourage everyone to check out both courses regardless of your experience level or familiarity with python or CAVE. Our courses can be accessed from [our elearning website](#).

NSF Unidata’s AWIPS group has the highest support interactions out of all the individual projects, as seen here in the [latest support reply summary](#). Through our support signatures we still have an active [support evaluation survey](#). The majority of our feedback has been overwhelmingly positive, and the graphic below is a summary from all responses we’ve received regarding the quality of service we provide:

Service Quality



Some of the latest open-ended feedback from the support evaluations includes the following:

- *Super fast response and solved issue. Thank you so much!*

- *You guys rock! Thanks for helping me switch to the AWIPS VM temporarily. :)*
- *Tiffany and Shay are always knowledgeable and helpful every time including yesterday.*
- *Incredible support and help, always just an email away. Very grateful for their work.*

The AWIPS team was able to send Tiffany to AMS 2024 where she hosted and led the [Unidata AWIPS Student Workshop](#), which was a great success and had an estimated 70 attendees. The workshop consisted of presentations on how AWIPS in the classroom can lead to important experience and qualifications when applying for jobs in the National Weather Service, demonstrations of CAVE and python-awips, and a first-hand telling from Victoria Elliott (a Texas A&M grad student) of how AWIPS is used in the classroom in their Meteorology programs.

In addition to the workshop, Tiffany was at our NSF Unidata booth during the student conference career fair, where we had the opportunity to interact with students of all grades from around the country. The team gave an [AWIPS update presentation](#) during the [second AWIPS System Updates session](#) later in the week.

The AWIPS team has had discussions with management as to if we're allowed to sell data "subscription" services to commercial and private sector entities. This idea may also include a paid EDEX installation workshop, as well as some possible EDEX maintenance workshops that the team could put on periodically and sell to our commercial sector users. The material developed for these workshops would be available to our academic community free of charge.

Lastly, since the last update, the AWIPS and Science Gateway teams have had to deal with significant Jestream2 (JS2) disruptions. Unfortunately, this presents to our users and EDEX outages and is bringing down our overall up time for our end users. The situation is complicated though, because the partnership with JS2 allows us to support and maintain many more machines than would be fiscally possible with any commercial cloud provider.

Software Releases

Since our last status report we have put out two production releases for [version 20](#). These updates include the proper inclusion of the gridslice package for both MacOS and Windows CAVE which allows the rendering of isentropic data, various CAVE and EDEX updates, and continued fixes and enhancements for each of the supported operating systems. Additionally, we upgraded our release of python-awips to [version 20](#) for both pip and Conda installation managers.

We successfully moved from supporting version 18 to version 20 before the end of 2023, and assisted our known EDEX maintainers at Texas A&M during the winter school break with their EDEX and CAVE migrations.

The team has since been actively developing a new major release of AWIPS – version 23.4.1. Significant progress has been made and so far we have been able to incorporate all previous Unidata-added functionality from version 20 into the new version 23. We are hoping to

release a beta of 23.4.1 by early May, and aim to have a production release for all three operating systems (MacOS, Windows, and Linux) before the CentOS7 EOL on June 30th. So far, we have also been able to create a beta version of python-awips for v23 and it seems to be functioning well in our existing Jupyter notebook examples.

Activities Ongoing/In-Progress

AWIPS development activities are constantly ongoing. Currently the following activities are in progress:

- The AWIPS team is actively developing our first beta release of v23 on Rocky 8, which is based on RHEL8 (instead of RHEL7)
- The AWIPS team has been maintaining our new Jetstream2 platform with multiple production and development EDEX servers for multiple different versions (v18, v20, and v23), with the help of the Science Gateways team.
- The AWIPS team is responding to all AWIPS support questions from the community and striving to provide realistic solutions in a timely manner.
- The AWIPS team is actively updating and refining our online documentation to be as accurate and useful as possible.
- The AWIPS team has maintained a bi-weekly blog series called AWIPS Tips that began on April 7th, 2021 and has been used to highlight useful functionality and fundamentals for CAVE, EDEX, python-awips, and general AWIPS announcements.

Future Activities

Future plans are constantly evolving to meet the needs of our users. The AWIPS team is focused on developing and releasing a production version of AWIPS on RHEL8 (Rocky 8) before the CentOS7 EOL. When this goes live, we will also update our python-awips package to have the latest v23 installation available on Conda and pip. We are actively participating in conferences, workshops, and virtual message boards (blogs) to expand our user base. Finally, the AWIPS team is testing the waters on developing possible business agreements with commercial and private entities to help bring in additional funds.

Metrics

Downloads November 1, 2023 - March 31, 2024

AWIPS downloads: 3,482

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Managing Geoscience Data**

The cloud-based EDEX data server continues to see widespread use and growing adoption. More and more datasets continue to be added to the server as Unidata deploys more decode/ingest nodes.

2. **Providing Useful Tools**

All AWIPS tools (EDEX, CAVE, and python-awips) are freely available, and also incorporate LDM/IDD technology for accessing geoscience data.

3. **Supporting People**

At this juncture, we are providing full technical support with regards to AWIPS for the community.

Prepared *April 2024*

Status Report: Science Gateway and Cloud Computing Activities

November 2023- April 2024

Shay Carter, Julien Chastang, Nicole Corbin, Ethan Davis, Doug Dirks, Tara Drwenski, Ana Espinoza, Ward Fisher, Thomas Martin, Ryan May, Tiffany Meyer, Jennifer Oxelson Ganter, Mike Schmidt, Tanya Vance, Jeff Weber

Executive Summary

- LROSE collaboration with CSU and NCAR
- WRF on Jetstream2 for classrooms and TCUs
- Development phases of Science Gateway Re-Imagined outlined
- Enhanced AWIPS EDEX servers on Jetstream2 Cloud
- JupyterHubs for atmospheric science education including GPU and Dask Hubs
- Maintained NEXRAD THREDDS server
- Updated and streamlined Docker image deployments for LDM, RAMADDA, THREDDS
- Strengthened cloud security
- Presented at Science Gateways Conference, 2023

Questions for Immediate Committee Feedback

1. What new or existing software are you excited to use or teach in your research? Does the installation process pose a barrier to entry, and can the Unidata Science Gateway team help you?
2. Are you interested in being a beta tester of the Re-Imagined Science Gateway?
3. Are there any specific case studies or success stories we should highlight?
4. What feedback do you have on the phased SGRI development approach?

Activities Since the Last Status Report

LROSE Collaboration between Colorado State University and NSF NCAR EOL

Unidata science gateway staff collaborates with Professor Michael Bell's team at Colorado State University and NCAR EOL to develop their science gateway, which features a JupyterHub equipped with LROSE (Lidar Radar Open Software Environment) radar meteorological software. Currently, we are adapting the SAMURAI (Spline Analysis at Mesoscale Utilizing Radar and Aircraft Instrumentation) software for JupyterHub use. We share our expertise in JupyterHub and related technologies with the team. Along with the LROSE team, we are in the beginning stages of drafting a submission to the ERAD (European

conference on RADar in meteorology and hydrology) conference and a potential submission to the Gateways conference. As part of this effort, Unidata has secured \$70K to fund Ana and Julien's contributions to the project.

Weather Research Forecast Model on Jetstream2

Summary

We continue to have the capability of running containerized versions of the Weather Prediction and Forecasting (WRF) numerical weather prediction model on Jetstream2. This capability has two major advantages: easy distribution of the model to any Docker enabled machine and the analysis and visualization of output in a data proximate manner, for example, in a JupyterLab environment.

WRF Navajo Technical University

Unidata is collaborating with the Southwestern Indian Polytechnic Institute and Navajo Technical University to deploy an operational WRF model over the Navajo Nation. This project aims to provide Tribal Nations, and the Tribal Colleges and Universities (TCUs) with the capacity for environmental monitoring in alignment with data sovereignty objectives. Progress in this area has been put on stand-by as the TCUs engage in other aspects of their project.

WRF Single Column Model in JupyterHub

In collaboration with Greg Blumberg at Millersville University, Unidata staff have deployed an idealized single-column WRF model in a JupyterHub environment for undergraduate instructional objectives. As a result of this collaboration, Unidata staff presented their procedures and findings at the Science Gateways 2023 Conference, hosted in Pittsburgh, PA on Oct 29 through Nov 1, 2023. Our collaboration with Dr. Blumberg is ongoing; he has requested the WRF SCM-capable JupyterHub to support course development and learning objectives for both the Fall 2023 and Spring 2024 semesters.

Unidata Science Gateway Re-Imagined

The Science Gateway Re-Imagined (SGRI) team—consisting of Julien Chastang, Nicole Corbin, and Ana Espinoza with managerial support from Ethan Davis and Tanya Vance—convenes bi-weekly to move the project forward. Our recent endeavors involve familiarizing ourselves with the Drupal Content Management System, the technology which UCAR plans to use to standardize its web presence. We are synchronizing our efforts with Doug Dirks and the Unidata web group to ensure our initiative moves in harmony with Unidata's overall web strategy.

Phased Approach to Development

The SGRI team, in collaboration with the Unidata Management Team, has developed a phased release approach to advance the project's development. These different phases, summarized below, were decided upon based on several factors, including ease of development, available

FTE's, and priorities which align with other Unidata goals. Most phases include an initial beta release. We anticipate soliciting participation from NSF Unidata committees and other community members to provide feedback on the design and usability of these features. Dates in parentheses are the tentative start/release dates for each phase of the project.

Phase 1: Requests and Education – *Users can request both compute resources (in the form of JupyterHubs) and educational resources (trainings, modules, etc.) and browse existing educational resources*

- a) (Aug 2024) JupyterHub Requests Beta
- b) (Oct 2024) Education Hub/Requests Beta
- c) (Jan 2025) Phase 1 Release

Phase 2: On-Demand Notebooks, Data Integration, and Community Hub – *Users can interact with Unidata curated “on-demand” notebooks without the need for a JupyterHub request, access data which is proximate to the computational environment, and share and develop ideas with colleagues in a community forum*

- a) (Jan 2026) Phase 2 Beta
- b) (Apr 2026) Phase 2 Release

Phase 3: Community Contributions – *Users can contribute to the content (educational materials, notebooks, workflows, etc.) found on the Science Gateway according to written guidelines for the management and maintenance of this content*

- a) (July 2026) Phase 3 Release

Phase 4: App Streaming & Fully Re-Imagined Science Gateway – *Users can “test-fire” Unidata products such as the IDV or Unidata’s version of AWIPS CAVE in their browser as a substitute for or prior to a local installation*

- a) (Jan 2027) Phase 4 Beta
- b) (Jul 2027) Phase 4 Release

JupyterHub Activities

Streamlining Launching of JupyterHub Clusters

We have deepened our understanding of Ansible and its capacity to enhance our JupyterHub deployment process, including distributing SSH keys and disabling non-essential Ubuntu services like CUPS for security. Consequently, we have integrated new Ansible playbooks into our existing deployment process to achieve these objectives.

GPU Cluster Developments

We collaborated with Alex Haberlie from Northern Illinois University and his master's student, Jeremy Corner, to set up a GPU-enabled JupyterHub. This platform supports their atmospheric science and machine learning objectives and is accessible at jupyterhub.unidata.ucar.edu. Access is granted upon request for those who wish to experiment with GPU technology.

We also collaborated with Maria Molina from University of Maryland to launch a GPU-equipped JupyterHub for her graduate level AI/ML class.

We have made progress in two key areas: optimizing our GPU-specific container images, which simplifies our GPU JupyterHub deployment and reduces disk image size; and developing a workflow that includes both CPU and GPU-enabled instances for improved Service Unit (SU) management. We are currently testing this hybrid CPU/GPU approach with our colleague Sean Freeman at the University of Alabama Huntsville.

Dask Cluster Developments

We worked with Professor Hanna Zanowski from the University of Wisconsin-Madison to establish a JupyterHub Kubernetes Dask cluster for her Earth System Science class of 15 students. Dask, a Python library, facilitates the scaling of Python code for parallel computing on distributed clusters.

JupyterHub Servers for Workshops, Fall and Spring Semesters

Unidata is employing our Jetstream2 resource allocation for the benefit of students in the atmospheric science community by providing access to customized JupyterHub servers at an accelerating pace. Unidata tailors these servers to the requirements of the instructors so they can accomplish their Earth Systems Science teaching objectives. Since the fall semester of 2023 (encompassing the length of this status report), 444 students at 14 academic institutions and various workshops have used Unidata JupyterHub servers running on Jetstream2.

Notably, we provided JupyterHub resources to:

- AMS 2024 Workshops
 - Student Python
 - Metpy
- Fall 2023 Machine Learning for Earth Systems Science two day workshop run by Thomas Martin and NCAR staff.
- Seoul National University students studying atmospheric science under an NCAR affiliate

Ongoing Activities

NOAA Big Data Program

- Unidata continues to manage the NEXRAD archive in Amazon S3, ensuring that realtime data are successfully delivered to the noaa-nexrad-level2 bucket. LDM is employed to deliver these data.
- Unidata also continues to deliver NEXRAD level 3 products to the unidata-nexrad-level3 bucket, part of the AWS public datasets program.

- TDS on Jetstream2 for level II NEXRAD: <http://thredds-aws.unidata.ucar.edu/thredds/catalog.html>
- AWS Explorer (Public access): <https://s3.amazonaws.com/noaa-nexrad-level2/index.html>
- Public Bucket for level II NEXRAD: <https://noaa-nexrad-level2.s3.amazonaws.com>
- Continue to populate the NEXRAD level II archive with real time data.
- Continue to populate new GFS 0.25 degree output and NCEP HRRR output to an S3 bucket for access. We did not place a TDS on this collection as this output is available from our standard sources.

University of Oklahoma REU Students

Unidata continues to collaborate with Ben Schenkel (OU) to provide data sets via the science gateway RAMADDA server. We also deployed a JupyterHub server so that NSF REU students at OU could access those data for their projects.

Andrea Zonca Collaboration

Unidata staff continues to collaborate with Andrea Zonca (SDSC/Jetstream2) employing his part of the "Zero to JupyterHub with Kubernetes" project to OpenStack and Jetstream2. We give Andrea feedback by testing his [instructional blog entries](#) and workflows. When we encounter issues, we submit bug reports via GitHub and work together until the problem is resolved. In addition, when we develop techniques or improvements to the workflow, we work with Andrea to ensure that this information is shared for the benefit of the wider Jetstream2 community.

Docker Containerization of Unidata Technology

We continue to employ Docker container technology to streamline building, deploying, and running Unidata technology offerings in cloud-based environments. Specifically, we are refining and improving Docker images for the LDM, RAMADDA, and THREDDDS. In addition, we also maintain a security-hardened Unidata Tomcat container inherited by the RAMADDA and THREDDDS containers. Independently, this Tomcat container has gained use in the geoscience community. To keep our containers up-to-date, especially with respect to security, we programmatically monitor and respond to upstream updates by automatically building and deploying the refreshed containers to DockerHub.

AWIPS EDEX in Jetstream2 Cloud

Unidata continues to host our publicly accessible EDEX servers on the Jetstream2 cloud platform where we serve real-time AWIPS data to CAVE clients and the python-awips users. The distributed architectural concepts of AWIPS allow us to scale EDEX in the cloud to account for the desired data feed (and size). We continue using Jetstream2 to develop cloud-deployable AWIPS instances as imaged virtual machines (VMI) available to users of OpenStack CLI. Since last summer all EDEX servers have been running on Jetstream2. Unfortunately, the service has not been entirely seamless and both the AWIPS team and the Science Gateway team have spent significant time troubleshooting and repairing machines to keep our servers operational. In addition, we have created custom CentOS 7 and Rocky 8

images for deployment on Jetstream2 on which to provision new EDEX machines before CentOS 7's End of Life on June 30, 2024. We have successfully created and launched a Rocky 8 EDEX system which the AWIPS team has been using to develop the latest version of AWIPS.

EDEX is designed so different components can be run across separate virtual machines (VMs) to improve efficiency and reduce latency. Our current design makes use of three VMs: one large instance to process most of the data and run all of the EDEX services including all requests, and two other ancillary machines which are smaller instances used to ingest and decode radar and satellite data individually.

We are currently supporting 4 sets of servers as described above: one set has been running our v18 software (as a backup during the transition from v18 to v20), two sets are running our new production v20, and a final system, using the Rocky 8 OS, has been used in the development of the new v23 release. Live backups allow us to be able to patch, maintain, and develop our servers while still having a fail-safe when something goes wrong with the current production system. We plan on decommissioning the v18 system in the near future, and will likely replace it with a second v23 server to use as the backup and development server, while the existing one can become the production v23 server. Some period after that we will decommission the "old" v20 servers after we have fully migrated to support v23.

Nexrad AWS THREDDS Server on Jetstream2 Cloud

As part of the NOAA Big Data Project, Unidata maintains a [THREDDS data server](#) on the Jetstream2 cloud serving Nexrad data from Amazon S3. This TDS server leverages Internet 2 high bandwidth capability for serving the radar data from Amazon S3 data holdings. TDS team member, Tara Drwenski, and Science gateway staff collaborate to maintain this server.

The URL for the THREDDS Nexrad radar server will be changed from thredds-aws.unidata.ucar.edu to tds-nexrad.scigw.unidata.ucar.edu to better reflect its purpose. It will also be integrated into the science gateway's data services section.

Jetstream2 and Science Gateway Security

We continually work with Unidata system administrator staff to ensure that our web-facing technologies and virtual machines on Jetstream2 adhere to the latest security standards. This effort involves such tasks as ensuring we are employing HTTPS, keeping cipher lists current, ensuring docker containers are up-to-date, limiting ssh access to systems, etc. It is a constantly evolving area that must be addressed frequently.

Unidata Science Gateway Website and GitHub Repository

Website

The [Unidata Science Gateway web site](#) is regularly updated to reflect the progress of what is available on the gateway. The news section is refreshed from time-to-time for announcements concerning the gateway. The conference section and bibliography is also

maintained with new information. We are in the process of redesigning this web site. See “Unidata Science Gateway Re-Imagined” section above.

Repository

All technical information on deploying and running Unidata Science Gateway technologies is documented in the [repository README](#). This document is constantly updated to reflect the current state of the gateway.

Presentations/Publications/Posters

- J. Chastang and A. Espinoza. Advancing Atmospheric Science Education: Customized PyAOS JupyterHubs via the Unidata Science Gateway. *In Proceedings, 40th Conference on Environmental Information Processing Technologies, 104th AMS Annual Meeting, Baltimore, Maryland, USA, 28 Jan.--1 Feb. 2024. American Meteorological Society.*

New Activities

Over the next three months, we plan to organize or take part in the following:

Forthcoming conference attendance

None planned at this time.

Jetstream2 Allocation Renewal

The ACCESS program has updated their procedures for requesting new or renewed “Maximize” allocations, the largest type of allocation grant. ACCESS begins accepting Maximize requests starting June 15, 2024 to July 31, 2024. If granted, Unidata’s Jetstream2 allocation will renew on October 1, 2024.

Over the next twelve months, we plan to organize or take part in the following:

Tomcat 8.5 End of Life

Tomcat 8.5 has reached end of life on 31 Mar 2024. This will require staff to transition the Tomcat Docker containers and any dependencies to the newer version of Tomcat.

Lessons Learned from April 14-17 Jetstream2 Outage

We aim to work with the Jetstream2 team to learn lessons from an outage that occurred from April 14-17.

Improved JupyterHub Kubernetes Cluster Stability

We aim to provide an optimal experience for our users, but unfortunately, we've experienced more downtimes than we'd prefer. Specifically, issues with disk attachments have disrupted users' ability to consistently access their Jupyter instances. To proactively address these issues, we plan to use cluster monitoring software like Prometheus and Grafana. This will allow us to identify and resolve problems before they impact the user experience.

Relevant Metrics

Fall 2023 / Spring 2024 JupyterHub Servers

Since spring of 2020, Unidata has provided access to JupyterHub scientific computing resources to about 1830 researchers, educators, and students (including a few NSF REU students) at 24 universities, workshops (regional, AMS, online), and the UCAR SOARS program. Below are the latest metrics since the last status report.

Fall 2023		
Boise State	0	Prof Lejo Flores
Florida Institute of Technology	8	Prof Milla Costa
Metropolitan State University of Denver	20	Erin Rhoades
Millersville University	3	Prof Greg Blumberg
University of Oklahoma	2	Ben Schenkel
University of Oklahoma 2	1	Professor Sakaeda
Southern Arkansas University	33	Keith Maull
University of Louisville	8	Prof Jason Naylor
University of Wisconsin	29	Pete Pokrandt, Prof Mayra Oyola
University of Wisconsin 2	20	Prof Hannah Zanowski
University of Wisconsin Dask	20	Prof Hannah Zanowski
CSU Python Workshop 1	25	Unidata Staff: Drew, Nicole, Thomas
CSU Python Workshop 2	14	Unidata Staff: Drew, Nicole, Thomas
AI/ML pyaos.unidata.ucar.edu November	36	Unidata Staff: Thomas
Colorado School of Mines	27	Prof Zane Jobe

Spring 2024		
AMS MetPy short course	27	Unidata Staff: Drew
AMS Python workshop	38	Unidata Staff: Drew
Florida Institute of Technology	14	Milla Costa
Florida State University	14	Christopher Homles
Millersville University	34	Greg Blumberg
Seoul National University	22	Duseong Jo
Southern Arkansas University	12	Keith Maul
University of Alabama Huntsville	12	Sean Freeman
University of Louisville	5	Jason Naylor
University of Maryland	17	Maria Molina
University of Northern Colorado	0	Wendilyn Flynn
University of Oklahoma	3	Ben Schenkel

Jetstream2 Allocation Usage Overview

In addition to service units (SUs) used for running various kinds of virtual machines—“regular” CPU and GPU instances—Unidata was also granted a limited number of compute, storage, and network resources to carry out Jetstream2 operations. These three kinds of resources are ephemeral, being created and destroyed as necessary. Thus, metrics regarding these resources are representative of short-term utilization, while SU usage is a metric that can be representative of our long-term Jetstream2 utilization.

Following Unidata’s 8M+ SU grant renewal, which went into effect October 2023, Unidata staff has been proactive in ensuring Jetstream2 resources are being used effectively in a non-wasteful manner with the on-going development of a SU monitoring script which tracks historical SU usage data and makes predictions on future usage.

SU usage and resource metrics, current as of April 17, 2024, are presented below.

SU Usage

Type	SUs Used	SUs Allocated	% Usage *
CPU	4,316,750	8,191,300	53 %
GPU	458,005	672,768	68 %

Resource Metrics

<u>Compute</u>			
Type	Used	Total	Percent Usage*
Instances	105	150	70 %
vCPUs	1178	4035	29 %
RAM	4.4 TB	15.8 TB	28 %

<u>Storage</u>			
Type	Used	Total	Percent Usage*
Volumes	222	400	56 %
Volume Snapshots	5	50	10 %
Volume Storage	32.4 TB	39.1 TB	83 %

<u>Network</u>			
Type	Used	Total	Percent Usage*
Floating IPs	52	310	17 %
Security Groups	73	100	73 %
Security Group Rules	229	300	76 %
Networks	3	100	3 %
Ports	134	250	54 %
Routers	2	15	13 %

* Percent Usage is rounded to the nearest whole number

Github Statistics*

Repository	Watches	Stars	Forks	Open Issues	Closed Issues	Open PRs	Closed PRs
science-gateway	6	17	11	5	167	10 (-4)	731 (+49)
tomcat-docker	11	64 (+4)	66 (+2)	0 (-2)	42 (+2)	0	88(+5)
thredds-docker	15	31 (+4)	27(+1)	3 (-1)	120(+3)	0	178 (+2)
ramadda-docker	4	0	2	1	10	0	35 (+1)
ldm-docker	9	12	14 (+1)	1	40	0	70(+5)
tdm-docker	5	4	7	0	10	0	24 (+1)

* Numbers in parentheses denote change from last stat report

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. Managing Geoscience Data

Unidata supplies a good portion of the data available on the IDD network to the Jetstream2 cloud via the LDM and the high bandwidth Internet 2 network. Those data are distributed to the TDS, ADDE, RAMADDA and AWIPS EDEX installations running on Jetstream2 for the benefit of the Unidata community. Unidata also makes the AWS Nexrad archive data accessible through the TDS Nexrad server running on Jetstream2 at no cost to the community. These data can be accessed in a data-proximate manner with a JupyterHub running on Jetstream2 for analysis and visualization. Containerization technology complements and enhances Unidata data server offerings such as the TDS and ADDE. Unidata experts install, configure and in some cases, security harden Unidata software in containers defined by Dockerfiles. In turn, these containers can be easily deployed on cloud computing VMs by Unidata staff or community members that may have access to cloud-computing resources.

2. Providing Useful Tools

Jupyter notebooks excel at interactive, exploratory scientific programming for researchers and their students. With their mixture of prose, equations, diagrams and interactive code examples, Jupyter notebooks are particularly effective in educational settings and for expository objectives. Their use is prevalent in many scientific disciplines including atmospheric science. JupyterHub enables specialists to deploy pre-configured Jupyter notebook servers typically in cloud computing environments. With JupyterHub, users login to arrive at their own notebook workspace where they can experiment and explore preloaded scientific notebooks or create new notebooks. The advantages of deploying a JupyterHub for the Unidata community are numerous. Users can develop and

run their analysis and visualization codes proximate to large data holdings which may be difficult and expensive to download. Moreover, JupyterHub prevents users from having to download and install complex software environments that can be onerous to configure properly. They can be pre-populated with notebook projects and the environments required to run them. These notebooks can be used for teaching or as templates for research and experimentation. In addition, a JupyterHub can be provisioned with computational resources not found in a desktop computing setting and leverage high speed networks for processing large datasets. JupyterHub servers can be accessed from any web browser-enabled device like laptops and tablets. In sum, they improve "time to science" by removing the complexity and tedium required to access and run a scientific programming environment.

3. Supporting People

A Unidata science gateway running in a cloud computing setting aims to assist the Unidata community arrive at scientific and teaching objectives quickly by supplying users with pre-configured computing environments and helping users avoid the complexities and tedium of managing scientific software. Science gateway offerings such as web-based Jupyter notebooks connected with co-located large data collections are particularly effective in workshop and classroom settings where students have sophisticated scientific computing environments available for immediate use. In the containerization arena, Unidata staff can quickly deploy Unidata technologies such as the THREDDS data server to support specific research projects for community members.

Status Report: Community Services

November 2023- April 2024

Nicole Corbin, Doug Dirks, Tanya Vance, Jeff Weber

Executive Summary

In addition to “normal, day-to-day” activities of communication and coordination with community members, the Community Services group’s efforts in the past six months have been focused on:

- Outreach and engagement with historically marginalized and underrecognized communities, including activities with tribal colleges and universities and the Sovereign Data Network project partners.
 - Some notable activities include live data streaming and data access from SIPI and NTU, collaborating with Lamont-Doherty Earth Observatory (LDEO) on the Hydrogeologic Impacts of Terrestrial Expressions of Mesozoic Perturbations (HI-TEMP) drilling project on the Navajo Nation, and working with The Haskell Foundation, Nebraska Indian Community College, Oglala Community College, Aaniiih Nakoda College, and Rising Voices Changing Coasts Hub to expand the Sovereign Data Network
- Outreach to the Earth System Science community through participation at the following:
 - Conferences: AMS, AGU, AIHEC, Rising Voices, and Environmental Data Science Innovation & Inclusion Lab (ESIIL) Innovation Summit
 - Workshops and Events: MSU Career Day, SciEd Career Highlights, the NSF NCAR Earth System Science and Technology hubs (NESTs) Community Workshop on Conceptualization and Coastal Resiliency, A New Lens for Community Engagement Workshop, and the Exploring Data Sovereignty and Sovereign Data Network Workshop
 - Working Groups: NSF NCAR Convergence Science Community Network, NSF NCAR-MSI Collaborators Community of Practice, NSF SOARS, NSF Earth Data Relations, and NSF ESIL
- Expansion of learning materials, resources, and offerings
- Progressing the Reimagined Science Gateway and integrated Education Hub to make our educational services more discoverable on the web as well as promoting the variety of educational services we are able to provide, including updating our [website](#) and community [requests](#)
- Development and delivery of the first module on Machine Learning Foundations in the Earth Systems Sciences at Metropolitan State University of Denver for [NSF Award #2319979](#)
- Facilitating a MetPy AMS Short Course, Python and AWIPS Student Conference sessions, and an AMS oral presentation on data readiness microlearning
- Partnering with COMET and USGS on the development of a series of NetCDF/CF microlearning resources and delivery of the Metpy for Quantitative Analysis training Spring - Summer 2024
- Providing coordination for the 2024 NSF Unidata Community Equipment Awards

- Supporting submission and post-submission actions of NSF Unidata’s next core funding award proposal
- Leading the planning and development of UCAR Community Programs (UCP)’s Strategic Plan and supporting the review of UCAR’s Strategic Plan

Questions for Immediate Committee Feedback

No questions at this time.

Activities Since the Last Status Report

News@Unidata blog

Posts to the News@Unidata blog appear regularly, but not on a specific schedule. Some highlights:

- [Unidata Staff at AGU Fall 2023 Meeting](#)
- [Call for Proposals: Unidata 2024 Community Equipment Awards](#)
- [New data-focused microlearning modules on Unidata eLearning](#)
- [Andrea Zonca and Jeremy Fischer Receive 2023 DeSouza Award](#)
- [Offer: NSF Unidata Science Gateway JupyterHub Resources Available for Spring 2024 Courses](#)
- [NSF Unidata Summer Student Internships Available!](#)
- [NSF Unidata: Where Did That Name Come From?](#)
- [NSF Unidata Staff at AMS 2024 Meeting](#)
- [AMS 2024 Conference Highlights from the NSF Unidata Staff](#)
- [2024 DeSouza Award Nominations](#)
- [NSF Unidata Committee Nominations Open](#)
- [Register Now for the 2024 Pythia Cook-off!](#)
- [Exploring Indigenous Data Sovereignty](#)
- Software release information
- A new [series of posts on AI/ML topics](#)
- Many AWIPS Tips and MetPy Mondays episodes
- Community meetings and other announcements

Dependencies, challenges, problems, and risks include:

- Finding community members willing to contribute stories (or story ideas) for the blog is an ongoing challenge. We’re starting to make progress working with committee members, but there is more to do.
- A sufficient number of committee member nominations and diverse and inclusive representation across ESS
- A sufficient and diverse pool of NSF Unidata Equipment Award proposals
- Broadening community participation across ESS

Community Outreach and Services

The community services group continues to actively reach out to and engage with NSF Unidata community members.

Progress has been made on the following:

- The community services team has led efforts to submit the core proposal for the next 5 years building on the strategic plan
- Continue to work with AIHEC, NEON, ESIIIL, and UCAR/NSF NCAR staff on NSF funded MSI engagement and NSF funded Sovereign Data Network project
- Convened [Exploring Data Sovereignty and the Sovereign Data Network Workshop](#) (funds provided by NSF NCAR CORE Award) at SIPI with 40+ individuals representing 8 TCUs, 5 R1 institutions, and 10 agencies and organizations resulting in commitment by new institutions to participate in the network, expanding NSF Unidata's community, and partnering on follow-up funding.
- Participating on new one year grant NSF#2220614 "The Indigenous Data Governance in Open Data Working Group"
 - Publication submitted to Nature
- Participating with members of Oglala Lakota College and ESIIIL focusing on data access, visualization, and analysis for Earth System Science; and efforts are underway to investigate potential pathways to foster complementary integration with the NSF Unidata-partnered, Sovereign Data Network
- Continue to engage with historically marginalized and underrecognized communities and institutions as part of NSF Unidata's outreach efforts to groups such as Rising Voices, SACNAS, and AIHEC
- Continue to serve on the CUAHSI HIS and DEI standing committees
- We continue to actively support the NSF NCAR/NSF SOARS program including serving on the 2024 Internship Selection Committee and Advisory Committee
- Participate, and be an ambassador for, the NSF NCAR-MSI Collaborators Community of Practice, NSF NCAR Convergence Science Community Network, and NSF NCAR EdEC's Engagement Group to support awareness of UCAR/UCP/NSF NCAR opportunities
- Engage and support the NSF NCAR|UCAR Intern Professional Development Workshop
- Engage with the Arctic Research Consortium of the US on multidisciplinary projects
- We continue to update NSF Unidata's social media channels (Facebook, Twitter, Google+)
- We continue to publish short videos/screencasts on the [Unidata YouTube channel](#).
- Represent NSF Unidata at the National Weather Service Partners events
- Actively participate in Super Science Saturday
- Support UCAR/NSF NCAR media services by responding to requests from media outlets

Dependencies, challenges, problems, and risks include:

- Engagement with NSF Unidata social media streams among community members is not particularly high
- Engaging with new Earth System Science communities that have different resources,

- capacities, and expectations
- Facilitating community adoption of new technological services (cloud, etc)

Learning Services

The community services group has expanded efforts to promote learning NSF Unidata products and workflows.

Progress has been made on the following:

- [NSF #2319979](#) “Machine Learning Foundations and Applications in the Earth Systems Sciences” (Nicole Corbin, PI and Thomas Martin, Co-PI) in collaboration with Dr. Keah Schuenemann (MSU Denver) and Dr. Karen Kortz (Community College of Rhode Island)
 - Module 1 pilot delivered to MSU Denver’s Synoptic Meteorology class on 4/25. Module evaluation to follow.
 - Module 1 is a no-code conceptual introduction to supervised machine learning, problem framing, and evaluation.
 - Results from the pilot will be presented at this summer’s [Earth Educators’ Rendezvous](#)
- Shared two microlearning modules at the 2024 AMS Education Conference
 - [Multidimensional Data Structures](#)
 - [Getting Started with Siphon and THREDDS](#)
- Prototyping the Science Gateway, Reimagined website, including a comprehensive Education Hub for all NSF Unidata’s educational offerings. See Cloud Computing Activities for more details.
- Partnering with COMET and USGS on development of a series of NetCDF/CF microlearning resources, to be completed Summer 2024
 - Will also deliver an off-the-shelf instructor-led training for USGS in early May 2024 on Metpy for Quantitative Analysis

Dependencies, challenges, problems, and risks include:

- Assignment and allocation of resources to develop new material

Community Equipment Awards

The NSF provides the NSF Unidata Program Center up to \$100k in equipment grant funds each year. In alignment with the NSF Unidata 2024 proposal, the Equipment Awards Program is designed to broaden participation and promote the use of NSF Unidata tools and systems (e.g., THREDDS, NetCDF, IDV, GIS connections) to support education and research on various aspects of climate studies (e.g., diagnostics, change and impacts), by providing grants to be used in the procurement of new computers and equipment including upgrades to existing classroom and laboratory equipment.

Progress has been made on the following:

The 2024 Equipment Awards solicitation closed on April 12 and the review panel is meeting to evaluate the proposals received on April 26, ahead of the spring Users Committee meeting.

See the [2024 Awards](#) page for details.

Dependencies, challenges, problems, and risks include:

- Increasing participation in the Equipment Awards program, especially by MSIs, smaller programs, and others who have not participated in the past.

Ongoing Activities

We plan to continue the following activities:

- Support for governing committee activities (convening and coordinating discussions, facilitating awards, etc.)
- Ongoing development of news articles and blogs for publication through News@Unidata
- Seeking partnerships to build and deliver community learning and development resources and continue to expand NSF Unidata's educational services
- Continue to engage with historically marginalized and underrecognized populations and institutions
- Engagement with other Earth System Science organizations, professional societies, and working groups
- Support the pursuit of funding and bringing greater public awareness to NSF Unidata
- Engage other UCAR/NSF NCAR divisions regarding NSF Unidata software use and collaborative projects
- Ongoing work to transition NSF Unidata's website to UCAR-mandated system
- Active participation in the Hydroshare Advisory Committee (CUAHSI)
- Continue working with AIHEC, NEON, ESILL, and UCAR/NSF NCAR staff on an NSF funded MSI engagement and Sovereign Data Network project

New Activities

Over the next three months, we plan to organize or take part in the following:

- Participate and present at the World Cafe for 2024 Rising Voices 12th Annual Workshop on Co-creating Research, Policy, Practice, and Action: The Rising Voices of Indigenous Peoples and Partners in Earth Systems Science
- Present poster at the Earth Educators' Rendezvous: Machine Learning Foundations and Applications in the Earth Systems Sciences
- 2024 Community Equipment Awards and supporting the DeSouza Awards and new NSF Unidata committees member selection
- Coordination of off-the-shelf learning experience at the NSF NCAR|UCAR Intern

- Professional Development Workshop Series
- USGS / COMET / NSF Unidata NetCDF/CF and Metpy for Quantitative Analysis training
- Submission of follow-up proposal for Sovereign Data Network project
- Completion of the UCP Strategic Plan

Over the next twelve months, we plan to organize or take part in the following:

- Delivery of NSF Unidata Education Hub web presence
- Delivery of prototypes of all three Machine Learning in the Earth Systems Sciences educational modules
- Engage other historically marginalized communities and institutions
- Collaborative efforts and expansion of the Sovereign Data Network with other TCU's
- Follow up proposals and collaborative activities
- Support and implementation of annual award tasks for year one

Beyond a one-year timeframe, we plan to organize or take part in the following:

- Continued maintenance of the reimagined NSF Unidata Science Gateway, including community contributions and adding resources to the catalog
- Champion CARE and FAIR principles

Relevant Metrics

Statistics from the Community pages on the NSF Unidata web site. Comparisons are made with statistics from the previous six-month period.

All community pages

Most recent six months:

- 39,397 unique pageviews (34,692 in previous period)
- 18.2% of total unique pageviews (16.7% in previous period)

Top community pages

1. All blog pages
32549 unique pageviews (26395 in previous period)
83% of total community pageviews (76% in previous period)
2. www.unidata.ucar.edu/community
3937 unique pageviews (1874 in previous period)
10% of total community pageviews (5% in previous period)
3. www.unidata.ucar.edu/about
1686 unique pageviews (1240 in previous period)

- 4% of total community pageviews (4% in previous period)
4. www.unidata.ucar.edu/events
886 unique pageviews (4830 in previous period)
2% of total community pageviews (14% in previous period)

Social media statistics, April 17, 2024

1. # of Twitter/X followers: 2056 (up from 2013)
2. # of Facebook followers: 912 (up from 908)
3. # of YouTube subscribers: 3798 (up from 3508)
4. # of LinkedIn followers: 185 (up from 145)

Unidata eLearning statistics, April 17, 2024

1. Total unique users: 300 (up from 162)
2. Enrolled users in Learn AWIPS CAVE: 215 (up from 138)
3. Enrolled users in Learn Python-AIPS: 23 (up from 18)
 - Note that microlearning courses are available without an account and are unable to be tracked

Strategic Focus Areas

We support the following goals described in NSF Unidata Strategic Plan:

1. **Managing Geoscience Data**
We monitor and collaborate with data sources to stay apprised of impending changes and to advocate for the needs of our user community.
2. **Providing Useful Tools**
We provide online learning infrastructure (via the elearning.unidata.ucar.edu site) support for a variety of learning objectives.
3. **Supporting People**
We provide user workshops, tutorials, and community workshops to help build supportive relationships between community members.

We coordinate with our governing committees to find ways to expand NSF Unidata's community participation. We use our website, electronic newsletters, and social media to keep community members informed about enhanced data services, software tools, and cyberinfrastructure.

We participate in UCAR/NSF NCAR and NSF projects for historically marginalized and underrecognized communities and institutions (SOARS, AIHEC, outreach to HBCUs). We provide services and tools to facilitate education and research in diverse communities. We work to broaden the NSF Unidata community by participating in student and professional conferences.

Prepared *April 2024*

Status Report: Data Services

November 2023- April 2024

Mike Zuranski

Executive Summary

Progress on ~~the RTSTATS Revamp~~ RTSTATS-NG has resumed. Deployment of an early release is underway and additional features are being developed.

Revising how we collect data metrics: Improving consistency and confidence in our reporting. A focus will be on our IDD and ADDE statistics collection process as a result of a recent analysis highlighting challenges & opportunities for improvement.

User support of course continues to be a primary focus. See the **Status Report: Support** for inquiry metrics.

Questions for Immediate Committee Feedback

None at this time.

Activities Since the last status report

RTSTATS-NG

That's right, it's got a name! **Real Time IDD Statistics - Next Generation**, or RTSTATS-NG for short. By the time you read this we should have a minimum viable product version online and will continue to make improvements and feature additions in the near term. The URL for RTSTATS-NG will be rtstatsdev.unidata.ucar.edu, and that will likely be where we host beta versions as development continues into production.

In an effort to replace the current operational version of RTSTATS as soon as possible, we will continue to use Daryl Herzmann's backend. Development on a new backend design was taking longer than expected, and so far no deficiencies have been found with his. We will note any areas on this design we wish to improve upon, and a decision on whether to move forward with a full redesign of a new backend will be made at a later date.

Once we are comfortable that RTSTATS-NG is ready for production, we will replace the current operational version at rtstats.unidata.ucar.edu with this; no ETA at this time.

Progress has been made on the following:

- Front-End:
 - Improved reliability
 - Added interactive Topology page
 - Reach minimum viable product
- Deployment:
 - We should have a version online by the time you read this
 - Development will continue
 - New features can be expected this summer
- Next priorities will include:
 - Ensure reliability of service & improve where needed
 - Add topology features
 - Add latency features
 - Add export tools
 - Add interactive geographic navigational tool

Dependencies, challenges, problems, and risks include:

- Time availability for continued development.

IDD Cataloging

No further development on this has taken place, however I try to keep a version running to help me with other work.

Over time a problem or two are revealed and I either work to address them as they happen or make note of them for later, but in this way at least keeping it running continues to provide benefits, one of which is stress testing.

Progress has been made on the following:

- Database Indexing has significantly sped up query times
- There is roughly a one month record (from one week) as a result of the above improvement.
- Fix occasional minor bugs

Dependencies, challenges, problems, and risks include:

Really no problems or challenges, work is just paused in favor of more urgent projects.

Status and Take-Aways Thus Far

No significant updates on IDD cataloging this cycle. Keeping a local version running has proved beneficial to my work and I am looking forward to when I can put significant development time into this project again. Unfortunately that is not likely to begin until after this summer at the earliest.

Ongoing Activities

We plan to continue the following activities:

- Work towards deployment of RSTATS-NG
- Explore new tools for UPC staff
- Explore new data sets
- Assist users in finding and using data
- Resume development on IDD Cataloging when possible
- Better understand and help manage geoscience data in our industries today

New Activities

We plan to begin work on the following activities:

- Developing new methods of data metrics collection
- Develop a strategy for LDM configuration and log centralization and management

Revising Metrics Collection Methods

Earlier this spring we were tasked with coming up with data volume metrics, many of which are in these status reports. The request came with short notice but fortunately that turned out to be a vital detail.

During these metrics collections, data volume numbers for IDD and ADDE services came up suspiciously short. Upon review we determined a number of factors were in play in those lesser numbers, both real and artificial. Case in point: the service we rely on for IDD volume metrics lost the first quarter of 2023's data so we had to extrapolate. Days after we collected our metrics this service provider underwent a service migration which inadvertently lost all of their prior data (at least when it comes to the pertinent metrics we poll from that service); had we not been asked to collect these metrics sooner we wouldn't have been able to.

It was also noticed that something may have been lost in the staff transitions of late. Tom

Yoksas was primarily responsible for collecting many of these metrics, and while other staff were informed on these methods it is possible that some details were missed.

Furthermore, as both staff and hardware cycle and transition it is crucial that we maintain as much consistency as possible. Otherwise we will run into situations similar to these far more frequently.

We have decided to embark on improving our methods of data metrics collection, and already have a few ideas:

- We are looking into building an internal reporting feature into the LDM. This will allow us to collect metrics similar to rtstats but more tailored to these purposes.
- Combine metrics from different sources. We are looking into ways we can combine these metrics such that they lend additional insight and/or improve confidence.
- Changing the logging format, frequency and cycling of metrics logging. ADDE logging had a weekly resolution, that is already down to daily.
- Investigate new methods for metrics collections.

Work on this has not yet begun in earnest. While we have entered the brainstorm phase, we do not have an ETA on when any new methods may be implemented.

LDM Configuration and Log File Management

Given the number of LDM installations we maintain and that many of them are nodes comprising a cluster, it is sometimes a challenge to traverse the various logs for intel, and a similar challenge exists when it comes to managing LDM configuration files such as ldmd.conf and any pqact.conf files. This disadvantage rose to the surface during the above metrics collection process; it was noticed that there were a number of opportunities for improvement when it came to managing these files across dozens of nodes.

The desired effects include, but are not limited to:

- Making it easier to find information across IDD cluster nodes pertaining to user requests, inaccessibility issues, upstream outages and other logging events
- Making it faster to respond to configuration change needs. At present delicate care must be taken when restarting IDD cluster nodes due to the number of downstream connections they support. This process could be automated
- Ensure consistency and redundancy across nodes & their configurations
- Automated scanning for opportunities for improvement from IDD cluster logs

We will be looking into different methods of managing remote file management better. At

this time it looks like the fastest avenue to getting what we need will be to develop some custom scripting to interact with our LDM nodes, but that development will take time. Efforts here will likely begin in the coming months.

Relevant Metrics

N/A at this time

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Managing Geoscience Data**

Gain a detailed understanding of what data is available, where and in what format. Convey that information to our users and partners. Work with data providers to help distribute and support their data sets. Support users and UCP staff in finding and working with data. Identify opportunities to reduce our data footprint, including the removal of duplicate data sets, LDM/IDD management.

2. **Providing Useful Tools**

RTSTATS-NG, IDD Cataloging, metrics collection, monitoring and alerting solutions.

3. **Supporting People**

User email support, training workshops, transcend UPC silos and support UPC staff.

Prepared April 2024

Status Report: Data Standards and Technical Outreach

November 2023- April 2024

Ethan Davis, Ward Fisher, Hailey Johnson, and Ryan May

Executive Summary

Unidata's netCDF teams continues to engage with the Zarr community on:

- 1) Zarr support in both the netCDF-C and netCDF-Java libraries;
- 2) the development of the Zarr version 3 specification; and
- 3) the development of the GeoZarr convention.

Unidata continues to be active in efforts to advance the Climate and Forecast (CF) Conventions for netCDF.

Unidata continues to be active in several international standards bodies and other communities focused on data and technology including the World Meteorological Organization (WMO), the Open Geospatial Consortium (OGC), and the Earth System Information Partners (ESIP).

Questions for Immediate Committee Feedback

No questions at this time.

Activities Since the Last Status Report

NCZarr/Zarr Specification Efforts

As part of implementing Zarr support in both the netCDF-C and -Java libraries, the NCZarr convention/extension has been developed to provide a clean and complete mapping between the netCDF and Zarr data models. During this work, the netCDF developers have been participating in discussions around clarification and evolution of the Zarr (version 2 and 3) specifications.

Progress has been made on the following:

- Members of Unidata's netCDF teams have membership on the Zarr Implementation Committee and the Zarr Enhancement Protocol (ZEP) Committee and regularly participate in the bi-weekly Zarr Community and ZEP calls.
- Members of Unidata's netCDF team participated in discussions around the GeoZarr convention. GeoZarr builds on the Climate and Forecast (CF) Conventions for netCDF and will be developed within the OGC community standards process.
- Members of Unidata's netCDF team are assisting with GeoZarr interoperability testing

of netCDF-C and -Java libraries.

CF Conventions for netCDF activities

Unidata has a long history of involvement in the development of the [Climate and Forecast \(CF\) Conventions for netCDF](#). These efforts continue with ongoing participation in development conversations on the [CF GitHub repositories](#), participation in and help in organizing the annual CF Workshops, and participation in the governance of CF.

Progress has been made on the following:

- The 2024 CF workshop ([announcement](#)) will be held at SMHI, Norrköping, Sweden on 17-20 September 2024.
- Ethan Davis continues serving as chair of the [CF Governance Panel](#).

OGC activities

Planning has begun for UCAR to host the Oct 2025 OGC Member Meeting.

The OGC netCDF Standards Working Group (SWG) met to discuss GeoZarr and its use of CF Conventions with members of the OGC GeoZarr SWG.

Ongoing Activities

We plan to continue the following activities:

- Track and engage in WMO data standards efforts
 - Ethan Davis is a member of the WMO Expert Team on Data Standards (ET-Data) and its Task Team for CF-netCDF (TT-CFNetCDF)
 - WMO CF-netCDF Profiles have been developed for radar data, oceanographic glider data, and aircraft data. Experimental distribution of CF-netCDF data on the WMO Information System (WIS) 2.0 is planned.
 - The WIS 2.0 provides similar functionality to the GTS as well as more interactive access to data. We plan to take a closer look at WIS 2.0 technologies to better understand possible connections with the IDD/LDM and how WIS 2.0 may impact and benefit the University community. We have had meetings with NWS and Environment Canada on how the LDM and THREDDS might interact with WIS 2.0 and its underlying OGC standards.
- Continue efforts to update and reorganize the NetCDF User's Guide (NUG)
 - Separate the aspect of netCDF that are useful to any user/developer, independent of which library or tool they use (i.e., data model, file formats, CDL definition, conventions, and best practices) from those that are library or language specific and
 - Clarify where and how the netCDF community can ask questions about the NUG as well as discuss and contribute to the development and advancement of the NUG.
- Continue efforts to registering netCDF Media Type (application/netcdf) with IANA

- The netCDF media type has been added to IANA's provisional registry list with Unidata listed as the standards-related body supporting the effort.
- Next Step: Complete documentation and metadata needed for full registration. The updated NUG will feed into this effort
- Represent Unidata in Earth System Information Partners
 - Unidata has been a Type II ESIP Partner Organization since 1999
 - Ethan Davis is currently the Unidata voting representative to ESIP.
- Represent UCAR and Unidata in OGC and various OGC working groups
 - Ethan Davis is the UCAR voting representative to the OGC Technical Committee, Kevin Sampson (NCAR/RAL, GIS group) is alternate voting representative.
 - Participate in OGC MetOcean Domain Working Group (DWG) meetings.
 - Ethan Davis is co-chair of the OGC netCDF Standards Working Group (SWG)

New Activities

Over the next three months, we plan to organize or take part in the following:

- Organize regular meetings of the OGC netCDF SWG to continue discussions on 1) the connections between GeoZarr and CF; 2) extending CF grid mapping / CRS capabilities to better support GIS/GeoTiff geolocation; 3) improving netCDF and CF support for Unicode in data and in variable names.
- Participate in the planning of the 2024 CF Workshop.
- Continue conversations with CF, WMO, and others on the development of mappings between CF Standard Names and GRIB/BUFR variable names.

Over the next twelve months, we plan to organize or take part in the following:

- Participate and help run the 2024 CF workshop ([announcement](#)) to be held at SMHI, Norrköping, Sweden on 17-20 September 2024.
- Deploy a draft version of the new, library independent, NetCDF User's Guide (NUG).
- Submit request for full registration of the netCDF media type with IANA

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. Managing Geoscience Data

Unidata's various data standards efforts contribute to important tools for data producers, especially those that design and develop new data products, and for those that develop software tools for data management, analysis, and visualization.

Status Report: GOES-East/West, NOAAPort and Other Satellite Imagery

November 2023- April 2024

Mike Schmidt, Mike Zuranski, Stonie Cooper

Executive Summary

NSF Unidata continues to operate satellite downlink facilities for the NOAAPort Satellite Broadcast Network (SBN) and GOES-East and GOES-West rebroadcast services on behalf of UCAR/NCAR and the NSF Unidata community. All received products are then provided via the Internet Data Distribution system (IDD) in various feeds and via remote access provided by AWIPS EDEX, McIDAS ADDE and THREDDS Data Servers.

Details on various efforts related to maintaining this capability are presented below.

Questions for Immediate Committee Feedback

None at this time.

Activities Since the Last Status Report

- Repurposed a 3.7 m mesh dish located at FL-2 from GOES-15 (GVAR) ingest to NOAAPort SBN ingest, and re-pointed the existing FL-2 NOAAPort dish at the new SBN satellite, Galaxy 31.

These moves were made for three reasons related to the NWS movement of the SBN from Galaxy 28 (located at 89W) to Galaxy 31 (located at 121W):

- The need for continued periodic GOES-15 (which formerly acted as GOES-West) operations was obviated by the successful launch, test out, and christening of GOES-18 as GOES-West, and the movement of GOES-17 to a parking orbit at 104.7W where it will serve as an on-orbit spare for GOES-East and GOES-West
- A second FL-2 mesh dish is configured to ingest GOES-14 data whenever it is returned to service either for checkout or if it assumes an active role as a spare for either GOES-East or GOES-West
- The line-of-sight from the existing NOAAPort dish on the south side of FL-2 to Galaxy 31 is impaired by large branches from trees outside of FL-3 and FL-2

Getting the worst offending tree removed or, at least, pruned enough to mitigate the problems being caused by the branches is unlikely.

- Work towards establishing an additional GOES-R downlink in a locations that has an unimpeded view of the southern sky continues

See Ongoing Activities for additional information.

- Establish a new NOAAPort downlink at the NSF-owned Marshall field site

An effort to establish a satellite downlink facility at the NCAR Marshall field site (just south of Boulder) has been slowed by an NSF moratorium on any ground penetrations until an environmental impact assessment (NEPA) has been completed. A non-penetrating ground mount will be installed in the Marshall compound, and a 3.8 m dish will be installed on the mount in the coming weeks. Following the satellite pad and dish installations, electronics needed to complete the downlink will be installed, and ingest testing will begin.

After the Marshall installation is complete, and assuming that high quality NOAAPort ingest can be achieved, and the interference of the trees at FL-2 can not be mitigated, the existing FL-2 NOAAPort solid dish will be converted to GRB downlink as it has an unobstructed view of the GOES-East orbital slot. This conversion would require that existing quad-shielded RG-11 coax be replaced by a dual run of LMR-400 coax from the dish to the 2nd floor computer room, and the LNB on the dish outside of the FL-2 cafeteria be moved to the dish being repurposed.

Ongoing Activities

We plan to continue the following activities:

- Participate in UW/SSEC's "fanout server" sharing of GOES-East/West data (redistribution of the GRB-200 UDP unicast stream over TCP) for GOES GRB products.

We are feeding from SSEC's GOES-East/West fanout servers, and they are feeding from the ingest machine that we operate. Sharing of the feed streams has allowed SSEC and Unidata to minimize effects of solar and terrestrial interference.

- Ingest GOES ReBroadcast (GRB) streams from GOES-East/West in real-time

As described in previous status reports, the 4.5 m dish located on the eastern satellite pad at the NCAR Mesa Lab has been used for GOES-West ingest since terrestrial interference (TI) was observed when pointing at GOES-East (GOES-16) in the fall of 2017. Replacement of power poles and lines to the south and downhill from the Mesa Lab prompted us to run a test to see if TI was still a problem when pointing at GOES-16. The results of the multi-day test convinced us that we could return the dish on the eastern satellite pad to GOES-East ingest, and proceed with installation of a dish on the western pad and use it for GOES-West ingest.

In the spring of 2022, we were given a 3.8 m satellite dish that was being excessed by a private company that was relocating their operations. This dish will be installed on the western satellite pad at the NCAR Mesa Lab. The running of dual coax cables from the western pad to the main Mesa Lab machine room has been completed, so the next step is the physical installation of this dish on the existing mounting pole

- Continue to distribute GOES-East and GOES-West data via the LDM/IDD and serve the data via the TDS, ADDE and EDEX

The volume of data available in the SATELLITE datastream can be seen in:

http://rtstats.unidata.ucar.edu/cgi-bin/rtstats/iddstats_vol_nc?SATELLITE+oliver.unidata.ucar.edu

Future Activities

IDD NIMAGE Datastream

The IDD **NIMAGE** feed was repurposed a few years ago from a feed that only contains satellite image products distributed in NOAAPort to one that can include value-added satellite products. The question for the committee is if there are other products that should be added to the **NIMAGE** feed?

VALUE-ADDED Products

We welcome contributions of additional value-added Level 2 satellite products by community members.

To date, Texas Tech University (Eric Bruning), CSU/CIRA, and NOAA's Vlab have provided value-added Level 2 products created from satellite image and lightning scans, and these have been distributed to the community in the NIMAGE IDD feed.

SSEC Collaboration

Continue working with SSEC on their *fanout* approach that insulates GRB ingestion from expected (e.g., twice per year solar interference periods; etc.) and unexpected (e.g., TI caused) service interruptions

Relevant Metrics

- The volume of GOES-East/West GRB products, 15 GB/hour average and 20 GB/hour maximum, is the second most voluminous feed on the IDD. This can be seen in the

real-time statistics listings from any of the accumulators for our toplevel IDD relay clusters under the SATELLITE feed. For instance :

https://rtstats.unidata.ucar.edu/cgi-bin/rtstats/rtstats_summary_volume?oliver.unidata.ucar.edu

- Feeding data to a slowly growing list of sites via the IDD:

We are distributing all or part of the GOES-East/West GRB products to groups within UCAR/NCAR (including EOL, RAL and ourselves), U.S. Universities, U.S. Government (e.g. NOAA and Military sites) and a number of international sites.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Managing Geoscience Data**

Providing TDS, ADDE and EDEX servers for GOES-East/West imagery and products benefits the greater community by providing access to real-time observations from the U.S. operational satellite constellation.

2. **Supporting People**

Providing access to data in real-time has been a fundamental Unidata activity since its inception. Continuing to provide data enables Unidata sites to focus on their educational and research activities.

Prepared *April 2024*

Status Report: Internet Data Distribution(IDD)

November 2023- April 2024

Mike Zuranski, Stonie Cooper, Mike Schmidt, Jeff Weber

Executive Summary

Unidata continues to support, update, and enhance the data available via the IDD for the benefit of research and education. Included but not limited to adding new data formats, bridging the knowledge gap in newly introduced data, and providing statistics of data flow and composition.

Questions for Immediate Committee Feedback

None at this time.

Activities Since the Last Status Report

Internet Data Distribution (IDD)

IDD data volumes continue to increase especially when new datasets are made available.

The following output is from a Linux-based data server that the UPC operates on behalf of the community, lead.unidata.ucar.edu:

```
20240418
```

```
Data Volume Summary for lead.unidata.ucar.edu
```

```
Maximum hourly volume 124359.462 M bytes/hour
```

```
Average hourly volume 63359.113 M bytes/hour
```

```
Average products per hour 392779 prods/hour
```

Feed	Average (M byte/hour)		Maximum (M byte/hour)	Products number/hour
SATELLITE	15264.235	[24.092%]	20758.908	6527.979
NEXRAD2	10499.148	[16.571%]	13769.130	109733.667
NIMAGE	8157.437	[12.875%]	12197.203	7478.208
NGRID	6992.673	[11.037%]	12000.162	66519.479
HDS	4516.974	[7.129%]	8186.544	30300.083
FNEXRAD	4488.991	[7.085%]	4962.380	8558.083
NEXRAD3	3814.658	[6.021%]	5045.800	91566.667
EXP	3055.116	[4.822%]	4558.471	2541.708

GEM	3047.831	[4.810%]	11342.866	5175.562
CONDUIT	1968.457	[3.107%]	64975.640	5847.688
UNIWISC	992.154	[1.566%]	1174.309	924.667
NOTHER	281.755	[0.445%]	766.130	59.542
FSL2	188.142	[0.297%]	682.392	1275.417
IDS DDPLUS	89.833	[0.142%]	114.865	55867.104
LIGHTNING	1.708	[0.003%]	3.947	402.667

Data Distribution:

IDD CONDUIT feed:

Beginning the week of 4/15 and lasting through at least 4/19 (the date of this writing) CONDUIT has been down at the source. This is the result of serious instability at the College Park Data Center and teams there have been working tirelessly to address them. In fact a Critical Weather Day was issued for this week to ensure NCEP et al. have all the resources they need at their disposal. There is no ETR at this time.

IDD FNEXRAD NIMAGE and UNIWISC feeds:

We continue to create the content for the FNEXRAD (NEXRAD Level III national composites), NIMAGE (GOES-East/West Level 2 images and products, fully reconstituted images from NOAAPort tiles and with broadcast headers and footers stripped off to leave "bare" netCDF4 files), and UNIWISC (select GOES-East/West images converted to McIDAS AREA format for use in legacy systems like GEMPAK) feeds.

Existing Data Distribution:

The data volume seen in the **SATELLITE** (which is known as **DIFAX** in LDM distributions prior to v6.13.6) listing above represents all products received in the GOES ReBroadcast (GRB) downlinks that we installed in UCAR (currently GOES-18 at the NCAR Mesa Lab and GOES-16 at UCAR Foothills Lab 2). The data volume seen in the **NIMAGE** entry represents GOES-East/West ABI Level 2 imagery that has been reconstituted by stitching together tiles that are distributed in NOAAPort and all other Level 2 products. In both cases, binary headers and footers that are added to products before distribution in NOAAPort have been stripped

off leaving “raw” netCDF4 files. The **UNIWISC** feed represents the volume of 3 select channels (0.64um VIS, 6.2um WV and 10.3um IR) for all coverages (CONUS, FullDisk, Mesoscale-1 and Mesoscale-2) of GOES-East/West image products that are in PNG compressed McIDAS AREA format that is suitable for use in GEMPAK, the IDV and McIDAS-V, McIDAS-X, and AWIPS.

Challenges, problems, and risks:

- Due to the nature of how LDM Feed Types are defined and used, it is often difficult to know where to insert new products. Even though we’ve already done this with NIMAGE, we want to refrain from adding anything to any NOAAPort feeds to avoid confusion and potential issues. We also know that at least several Feed Types are currently going unused going out from the IDD, yet we cannot tell how sites use Feed Types if they do not report to RTSTATS. We don’t want to repurpose an existing Feed Type unless we can be sure it is and will be unused for its original purpose (e.g. DIFAX to SATELLITE). We are continuing to evaluate options to repurpose existing Feed Types, as well as looking forward to new options for Feed Type handling with future versions of the LDM. At this time, however, it can be tricky to find where the best place is for new products to go.
- Plans are also underway to enhance how we manage our LDM installations and IDD node configurations, which will make them more robust, easier to manage and faster to investigate & respond to issues. These activities have been paused due to higher priority projects and staffs’ limited time availability.
- The operational RTSTATS web site will be going away in the near future. The environment currently hosting this will be decommissioned soon, and previous attempts by NSF Unidata staff to port the old software to newer environments have not been successful. Mike Zuranski has continued his work on his replacement, RTSTATS-NG, and we hope to have this version online before the end of May (if it isn’t already by the time you read this). However, due to multiple learning curves with the chosen tech stack (Dash) and other priorities, this has taken longer than planned. It’ll be worth it though, I swear!
- We recently came to the conclusion we need to find better ways to collect metrics for IDD usage, as well as some of our other services such as ADDE. The Netvizura service we have been using to collect the metrics below had lost the first quarter of 2023’s data, and days after pulling these metrics they had a bad database migration that lost all metrics up to that point. Between that and working to better maintain consistency in these metrics we will be working to improve our methods. More details on this can be found in the Data Services report.

Ongoing Activities

We plan to continue the following activities:

- Fire weather products (HRRR Smoke) that are being made available by NOAA/GSL in an EXP feed were added to the set of HRRR products that are available from hrrr.unidata.ucar.edu. These products along with other model output are available via the TDS and Unidata AWIPS EDEX.
- Other data sets we continue to explore with NOAA/GSD/ESRL are:
 - [RRFS](#)

NOAAPort Data Ingest

- Ingest of the DVBS-2 NOAAPort Satellite Broadcast Network (SBN) products and their relay to end-users via the IDD has been “operational” at the UPC since August 2014.

Considerable effort has been expended in streamlining our NOAAPort ingest systems and assisting sites (UWisc/SSEC, NOAA/GSL, NOAA/SPC, Fox13 TV) in troubleshooting problems being experienced in their systems.

- The NOAAPort-derived data streams (**HDS, IDS|DDPLUS, NGRID, NIMAGE, NEXRAD3** and **NOTHER**) are redundantly injected into the IDD at four geographically separate locations: UCAR/Unidata, UWisc/SSEC, Allisonhouse.com and Fox13 TV in Tampa, FL.
 - LSU/Climate is no longer participating in this due to satellite dish hardware degradation and after consideration determined it was not worth the cost to invest in a new dish, installation and other costs.
- Unidata's NOAAPort ingest package is bundled with current versions of the LDM. The current LDM release is v6.15.0.

Relevant Metrics

- Approximately **545** machines at **176** sites are running LDM-6 **and** reporting real-time statistics to the UPC.

We routinely observe that the number of sites reporting real-time statistics fluctuates. We are not 100% certain why this may be the case, but our best guess is that some sites do not keep their LDMs running all of the time; campus firewall adjustments block the sending of the statistics; and/or sites decide to stop sending statistics. The latter possibility seems to be happening more frequently.


NB: We know that there are a number of sites that are participating in the IDD, but are not reporting real-time statistics back to us. Reporting of real-time statistics is not and never has been mandatory.

Unidata staff routinely assist in the installation and tuning of LDM-6 at user sites as a

community service. We have learned about sites not sending real-time statistics during these kinds of support activities, and a number of times the impediment to sending in stats is firewall configurations at the user sites.

- A number of organizations/projects continue to use the LDM to move substantial amounts of data that do not report statistics to Unidata: NOAA, NASA, USGS, USACE, Governments of Spain, South Korea, private companies, etc.).
- UCAR IDD toplevel relay clusters, **idd.unidata.ucar.edu** and **iddb.unidata.ucar.edu**

The IDD relay clusters, described in the June 2005 CommunitE-letter article Unidata's IDD Cluster, routinely relays data to more than 1250 downstream connections. The primary IDD relay cluster, **idd.unidata.ucar.edu**, was moved to the NCAR/Wyoming Super Computing facility in Cheyenne, WY in late August 2019.



Over the period from March 23, 2023 through December 31, 2023 (IDD volume snapshots are taken during periods that do not have monitoring dropouts in NetVizura plots) the average volume of LDM/IDD data flowing through the Front Range GigaPop averaged around 6.4 Gbps (~69.12 TB/day), and peak rates reached 9.3 Gbps (which would be ~100TB/day if the rate was sustained (which it is definitely **not**)).

The following table of volume snapshots shows that the volume of data flowing to downstreams out of UCAR has been reasonably consistent:

Date range	Src		Dst		Total	
	Ave	Max	Ave	Max	Ave	Max
20200508 - 20200630	5.4	7.5	42.1	52.9	5.5	7.5

20200701 - 20200930	5.4 7.9	41.9 60.3	5.4 7.9
20201001 - 20201231	5.2 6.9	39.9 55.9	5.3 7.0
20210101 - 20210331	5.5 8.0	42.3 59.9	5.5 8.1
20210401 - 20210415	6.1 15.5	46.4 112.7	6.1 15.7
20210601 - 20210719	6.6 9.2	50.5 73.0	6.6 9.2
20210908 - 20211005	7.6 14.9	59.3 121.7	7.7 15.0
20211101 - 20211231	6.7 9.1	52.4 71.4	6.8 9.2
20220208 - 20220311	6.6 15.2	53.5 114.8	6.6 15.3
20220412 - 20220521	7.2 14.5	52.6 103.7	7.3 14.6
20220717 - 20220831	7.3 13.3	46.3 86.1	7.3 13.4
20220714 - 20230313	7.8 11.7	51.1 77.4	7.8 11.7
20230910 - 20231013	6.8 11.7	39.4 74.3	6.8 11.8
20230323 - 20231231	6.3 9.2	37.0 56.5	6.4 9.3

NB: The units for Src and Total Ave and Max are Gbps (gigabits per second), and the units for Dst are Mbps (megabits per second).

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Managing Geoscience Data**

The IDD project demonstrates how sites can employ the LDM to move and process data in their own environments.

2. **Providing Useful Tools**

The freely available LDM software and the IDD project that is built on top of the LDM have served as a demonstration for distribution of real-time data for a variety of organizations including the U.S. National Weather service.

The cluster approach for LDM/IDD data relay that Unidata pioneered has been adopted by several Unidata university sites, and is currently being implemented at U.S. government sites.

Unidata's NOAAPort ingest package, which is bundled with LDM-6, is being used by a variety of university, U.S. government, and private sector entities.

Both the LDM and NOAAPort ingest packages are bundled with AWIPS.

3. **Supporting People**

The IDD is the primary method that core Unidata sites use to get the meteorological data that they need. Providing access to data in near real-time is a fundamental Unidata activity. The IDD-Brasil, the South American peer of the North American IDD, and IDD-Caribe, the Central American peer of the North American IDD, are helping to extend real-time data delivery throughout the Americas

Prepared *April 2024*

Status Report: IDV with RAMADDA

November 2023- April 2024

Yuan Ho, Julien Chastang

Executive Summary

We continue to support, update, and enhance the 3D data visualization and analysis tool IDV for our community. Our current activities include: coordinating with netCDF-Java group to add new data formats, collaborating with the SSEC developers to enhance the VisAD library, and working with our community to promote the usage of the IDV in research and education.

Questions for Immediate Committee Feedback

We have noticed that many advanced features of the IDV, such as formulas and trajectory displays, have not been widely used in the community and many data servers that the IDV can directly access are less well known to IDV users. We would like to provide help to classes, research groups and project teams to use these resources. Can committee members help to establish such connections?

Activities Since the Last Status Report

IDV Releases

The IDV 6.2u1 was released in October 2023.

The IDV 6.2u2 was released in December 2023.

IDV System Changes

__IDV Certificates__

Java Windows app and MacOS certificates have been renewed and will be valid until at least May 30, 2021 (MacOS certificate is valid until 2024). Moreover, as properly signing the IDV under these different environments can be an involved process, this information has been thoroughly [documented here](#).

__Changes to nightly release that will eventually be incorporated into into stable version__

- IDV uses the latest Java 8 AdoptOpenJDK

- IDV employs latest Java3D (1.6.2)
- Updated the IDV code signing certificates on all platforms (i.e., MacOS, Windows, Webstart)
- IDV now “notartized” on MacOS
- Updated Unidata's Install4J license from version 5 to 8.
- Updated the IDV Install4J configuration.

IDV Display Changes

__Pre/Post process data for ML applications __

The IDV offers a range of statistical analysis formulas, encompassing area averages, level averages, maximum, minimum, mean, percentiles, and summations. The results of these analyses can now be produced as non-geolocated data and exported in formats such as CSV or netCDF. This newly introduced feature empowers users to leverage the IDV's versatile access to multiple data servers, enabling them to preprocess data for applications including machine learning and other scientific uses. I am planning to further enhance the APIs and present these features in the coming AGU conference.

__ Multi-threading bug affecting the authentication process __

Addressing a critical aspect of the IDV system, I have dedicated efforts to resolve a multi-threading bug affecting the authentication process. Recognizing the significance of a seamless and secure authentication mechanism, I have undertaken a comprehensive review and debugging process to eliminate any potential issues arising from concurrent threads. By identifying and resolving this bug, I aim to fortify the overall stability and efficiency of the IDV authentication system, providing users with a smooth and secure interaction with the platform. This commitment to quality and responsiveness reflects our dedication to delivering a top-tier user experience within the IDV framework.

__IDV freezing and hanging bug fixed__

It appears that there is an issue when launching IDV in conjunction with the Grammarly desktop application, resulting in freezing and hanging. The specific cause of this behavior may be related to the interaction between the two programs. This problem was getting very serious with many IDV users reporting this problem. After exploring this problem, we reach the conclusion that conflicts of these two applications were at the system library levels and we have to wait for the update of Java 3d library. We suggest users to consider disabling Grammarly temporarily when launching IDV or set the focus to a third application right after launching the IDV to help mitigate compatibility issues.

__IDV Zoom Enhancement__

One persistent issue with the IDV has been the vanishing of Java 3D objects when zooming in to certain levels. This occurs when the clipping distances are not properly configured, causing objects beyond the specified range to go unrendered and disappear. Adjusting the clipping distances in Java 3D is essential to ensure the visibility and accurate representation of objects within the scene. Through troubleshooting and refinement of these parameters, coupled with the implementation of a new algorithm to delineate the relationships among the camera viewpoint, object positions, and clipping planes, we have successfully addressed the problem of disappearing Java 3D objects in the IDV. Users can now visualize the 3D display even at street levels, a pivotal feature for the hydro community in displaying and analyzing flooding at the street scale.

IDV Community Support

In the hybrid environment of in person and remote-learning system as a result of the COVID-19 pandemic, we keep helping universities and research institutes to run Unidata's Integrated Data Viewer (IDV) at home.

- Yuan worked with Jeff Weber to provide support to the NSF funded project "The Indigenous Data Governance in Open Data Working Group".
- On November 7, 2023, I had a zoom meeting with members of the Association of Consulting Meteorologists in a presentation that focused on the forefront of meteorological advancements. This presentation aims to illuminate the latest developments, with a particular focus on innovative tools and methodologies that can empower consulting meteorologists in their crucial roles. One of the members emailed the organizer to express that I was the best presenter that ACM has offered in this forum. This positive reception fuels my excitement to engage with the association, fostering an environment of shared knowledge and collaborative exploration.

MSU IDV Project

I have collaborated with professor Sun from MSU and submitted the proposal, "Scientific Visualization and Mathematical Modeling of Weather Data: An Interdisciplinary Approach to Learning with IDV (Interdisciplinary IDV)," to NSF. This joint endeavor seeks to weave together the realms of scientific visualization and mathematical modeling, using the Integrated Data Viewer (IDV) as a central tool. By capitalizing on the unique features of the IDV, we aim to provide a comprehensive platform for hands-on exploration of weather data, empowering learners to engage deeply with the intricacies of mathematical modeling in the context of atmospheric sciences.

This collaboration with MSU exemplifies a commitment to cross-disciplinary education and research, promising to contribute valuable insights to both meteorology and education. Together, we anticipate achieving impactful outcomes that advance the understanding and

application of scientific principles in the dynamic field of weather data analysis.

IDV Publication Highlights

[Synoptic-Dynamic Meteorology in 3D: Introducing an IDV-Based Lab Manual](#) by Gary Lackmann, B. Mapes and K. Tyle

A [Google Scholar Search](#) reveals a number of publications that cite use of the IDV ([doi:10.5065/D6RN35XM](https://doi.org/10.5065/D6RN35XM)).

IDV and RAMADDA Training, Conference Attendance and Presence

__2023 AGU Fall Meeting__

- Pre and Post-processing of Machine Learning Datasets with UNIDATA's IDV

__2024 AMS Annual Meeting__

- Three Dimensional Gridded Visualization and Analysis of individual NEXRAD Level2 Volume Radar Data in Unidata's IDV

Ongoing Activities

We plan to continue the following activities:

__Investigation of Java 3D Alternative__

Because of concerns about the long-term viability of the open-source Java 3D project, the IDV team has begun discussions with our University of Wisconsin, SSEC collaborators to replace Java 3D with a more viable alternative within the VisAD API. We have started investigating whether the [Ardor 3D](#) can meet that objective. Looking into alternatives to Java 3D was also a goal described in the [Unidata 2018 Five-year plan](#).

New Activities

Over the past few months, we plan to organize or take part in the following:

We plan to upgrade the version of OPenJDK Java. This change will necessitate in depth testings and the IDV building and distribution workflow.

Relevant Metrics

__E-Support__

The IDV team continues to provide the geoscience community with high-quality support through e-support software and idv-users mail list. In the last half year the IDV team has closed ~40 e-support tickets. Each individual ticket may and often does involve many back-and-forth messages. There is an especially large number of support requests coming from international users.

Top ten universities running IDV are: Millersville, Oklahoma, University of Utah, St Cloud state, Plymouth, NC State, West Kentucky, Lyndon State, University of Illinois, and San Francisco State.

__GitHub Pull Requests__

In the area of greater collaborative development, since the migration of the IDV project to github, we have closed a total of 125 “pull requests” or code contributions from internal and external collaborators.

__Youtube IDV Instructional Videos__

In the area of online IDV training, the Youtube IDV instructional videos have been viewed thousands of times.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Managing Geoscience Data**

The IDV is a state of the art geoscience visualization application. It gives users the ability to view and analyze a rich set of geoscience data, including real time data, in a seamless and integrated fashion. This analysis is captured in IDV bundles. RAMADDA is a content management system and service specifically tailored towards the sharing and distribution of IDV bundles facilitating distribution of scientific data and analysis.

2. **Providing Useful Tools**

The IDV has been an open-source project for several years. The IDV is available on the github version control platform for greater open-source collaboration. The IDV provides users the unparalleled ability to analyze, integrate, and visualize heterogeneous geoscience data in two, three, and four dimensions. The IDV coupled with RAMADDA enables geoscience specialists the capability to share and collaborate

their IDV analysis via social scientific networks.

3. **Supporting People**

Unidata offers yearly multi-day training and occasionally regional workshops for IDV and RAMADDA. The IDV coupled with RAMADDA enables our earth science community partners to distribute geoscience data and metadata through web-based technologies thereby fostering scientific collaborations. Moreover, the IDV's ability to share bundles through RAMADDA creates a scientific social and collaborative network for the geoscience community.

Prepared *April 2024*

Status Report: Information Technology

November 2023- April 2024

Mike Schmidt, Matt Perna, & Jennifer Oxelson

Executive Summary

Our role is to maintain and enhance the productivity of the staff and assist with the resolution of issues in service to the community. Primarily, that consists of keeping end-user and developer systems secure, and keeping servers and services highly available, patched, and operational for the community. This report is informational and there are no pressing issues.

Questions for Immediate Committee Feedback

Unless committee members or the community are experiencing performance issues that we could help resolve, no other feedback is requested.

Major Activities

- Unidata staff have been installing/vetting new support ticketing systems to replace the currently used eSupport package.
- UCAR continues to migrate select services to a centralized model and Unidata is involved in implementing the service(s) for our systems and users. Recently completed and/or anticipated are;
 - centralized backups (CrashPlan, complete)
 - centralized end-user security agent (Palo Alto Cortex XDR, complete)
 - staff net deployment (all office networks, complete)
 - corporate device management (in progress)
 - centralized access and identity management (in progress)
 - All systems have Disk Encryption enabled (Mac and Windows - Complete)
 - Working with help@ucar.edu to obtain systems for our staff to migrate to, and awaiting a timeline on this from them.

As of writing this, we have implemented 2 MDM solutions to manage our Macbooks as well as our remote Windows clients. Firewalls, antivirus, and overall configurations are monitored constantly and report back when there is an anomaly on both Operating systems. Cloud backups have been implemented on all Unidata client workstations for the past year and we see a foreseeable upgrade in the way we back up clients to the cloud in the coming year. Disk encryption will be enforced across the board on all client operating systems. Unidata IT can remotely manage any Unidata client workstation if it is connected to the internet.

UCAR's Mesa Lab Data Center (MLDC) co-location facility upgrade is mostly complete, and Unidata is scheduled to begin moving equipment to the MLDC during the work week of April 15th. There may be the need to roll some of our Internet-based services to the NWSC in

Cheyenne to avoid community visible service outages.

Daily, we continue efforts to keep services and systems secure which takes consistent attention and occasional herculean efforts (to patch everything all at once). UCAR continues to embark on new initiatives to segment the network into smaller and smaller zones and gain a more dynamic inventory of assets on the network. Unidata continues to play a role in these efforts.

Ongoing Activities

We plan to continue the following activities:

- Day-to-day system and network support to the community as needed
- Resolve daily staff help desk issues
- Maintain security profile and exceed UCAR security standards
- Following UCAR directives regarding cybersecurity initiatives

Prepared *April 2024*

Status Report: LDM

November 2023- April 2024

Stonie Cooper

Executive Summary

Unidata's LDM team continues to update source code and operating paradigms with ever-changing user demographics and user requirements, particularly in the area of security and inclusiveness of data.

Questions for Immediate Committee Feedback

As intimate stakeholders in data distribution and access for research and education, please let us know about any data distribution software needs that you personally experience that are not currently fulfilled by use of LDM.

Activities Since the Last Status Report

The LDM is the primary software package by which research and education institutions obtain near real-time meteorological and related data.

Progress has been made on the following:

- Template created for packaging precompiled self-extracting binaries for key Linux distributions
- Transition from previous lead developer to new lead developer
- Examination and testing of source code language to enhance system security
- Moving hard-coded aspects of LDM into configuration files.
- Answered many questions from Universities, NOAA, US Military, and corporations.

Dependencies, challenges, problems, and risks include:

Ever changing landscape of new data types, adjustments from data source paradigms.

Response to security concerns will impact installation paradigms.

Ongoing Activities

We plan to continue the following activities:

- Support and maintain the LDM
- Convert the LDM tangential applications to more secure computing language as an

- introduction to new coding paradigms.
- Continued design and testing of the next generation of data distribution mechanisms.
- Introduction of data bandwidth metrics to provide real-time data usage accounting.
- Migration of LDM training to on-demand video training sessions.

Relevant Metrics

- The LDM system at the Unidata Program Center powers the Unidata IDD (Internet Data Distribution) system. Metrics on that program can be found in the IDD status report.
- Number of LDM package downloads.
- Number of support and training requests.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Managing Geoscience Data**
By enabling researchers, teachers, and students to process a wide variety of meteorological and related data in near real time.
2. **Providing Useful Tools**
By enabling researchers, teachers, and students to obtain a wide variety of meteorological and related data in near real time and at no cost via the Internet.
3. **Supporting People**
By answering support questions, writing documentation, and conducting workshops.

Prepared *April 2024*

Status Report: netCDF

November 2023- April 2024

Ward Fisher , Hailey Johnson , Ethan Davis

Executive Summary

The netCDF team continues to work towards maintaining the reliability of the netCDF libraries, while keeping one eye forward as to the future needs of our community. We have continued our community engagement efforts and collaborations whenever and wherever possible; examples of this include our involvement with the Zarr Community meetings and our membership on the Zarr Enhancement Protocol (ZEP) committee. We have also continued conversation with the HDF group and other community groups working on similar efforts.

We continue to address the issues associated with the proliferation of new mainstream architectures (Apple Developed M1/2/3/ARM), evolving compilers and standards, and extending our collaborations with tangential, but related, projects (conda-forge libnetcdf feedstock, for example).

In the past several months, the NetCDF team has participated in a developer exchange program with the Atmospheric Chemistry Observations & Modeling (ACOM) at NCAR. The ACOM developer, Kyle Shores, was able to significantly modernize the build infrastructure for netCDF-C, freeing up resources for the core development team to work on ncZarr and S3 support, as well as user support and general quality-of-life technical improvements.

Questions for Immediate Committee Feedback

No questions at this time.

Activities Since the Last Status Report

Snapshot of NetCDF Development Status

We are using GitHub tools for C, Fortran and C++ interfaces to provide transparent feature development, handle performance issues, fix bugs, deploy new releases and to collaborate with other developers. Additionally, we are using docker technology to run netCDF-C, Fortran and C++ regression and continuous integration tests. We currently have **276** open issues for netCDF-C, **105** open issues for netCDF-Fortran, and **55** open issues for netCDF-C++. The netCDF Java interface is maintained by the Unidata CDM/TDS group and we collaborate with external developers to maintain the netCDF Python interface.

In the netCDF group, progress has been made in the following areas

since the last status report:

- Support for Amazon S3 access via libnetcdf (using either the Amazon S3 SDK library, or an internal interface layer) has been further improved.
- The netCDF and netCDF-Java teams continue to participate in the Zarr Community meetings, in order to help guide the development of the Zarr v3 and future specifications in a way that promotes broad compatibility across Zarr implementations.
- The netCDF and netCDF-Java teams have also joined with the Zarr Enhancement Protocol (ZEP) committee, in an effort to help codify the process by which features are added to the Zarr v3 specification.
- Continuing improvement for the NUG: We previously migrated the NetCDF User's Guide to a new, separate repository. This repository will contain the concise, language-agnostic summary of the netCDF data model. Language-specific documentation (primarily used by developers) will remain associated with the individual code repositories.
- Further enhancements to the netCDF-C documentation, modernization of the netCDF-Fortran and netCDF-C++ documentation.
- We continue to see a high volume of contributions to the netCDF code base(s) from our community, for which we are grateful. While these contributions require careful review and consideration, it is encouraging to see this model of development (enabled by our move to GitHub) being more fully embraced by our community.
- Improvement and collaboration on additional filter and plugin support for dynamic, selective compression, based on work contributed by Charlie Zender and Ed Hartnett.
- As a result of increased interest, the DAP4 functionality has been significantly improved. A corresponding set of changes was propagated to the NetCDF-Java code base. Some discrepancies in the DAP4 specification were discovered, and resolution is on-going.

Dependencies, challenges, problems and risks include:

- The increasingly small group of netCDF developers is under a lot of pressure to provide project management as well as implement new features, fix bugs, provide support, etc. With 1.5 FTE assigned to the project, the workload is significant.
- Difficult issues which require intense debugging can bog down progress in other areas of netCDF and related projects.
- Rapid evolution of the Zarr standard is very useful, but also provides a bit of a moving target.
- Increase in external contributions has greatly increased the project management overhead for netCDF-C/C++/Fortran.
- Advances in compilers (GCC 10.x) and newer architectures (such as Apple's ARM M1/M2 architecture) are requiring additional overhead to ensure compatibility.
- The proliferation of cloud environments requires specific attention.

Ongoing Activities

We plan to continue the following activities:

- Continue work towards adoption of additional storage options, separating out the data model from the data storage format (as much as possible).
- Improve the messaging around the expanded functionality of netCDF.
- Provide support to a large worldwide community of netCDF developers and users.
- Continue development, maintenance, and testing of source code for multiple language libraries and generic netCDF utility programs.
- Continue modernizing the documentation for netCDF-C, Fortran and C++ libraries.
- Extend collaboration as opportunities arise, for increasing the efficiency of parallel netcdf-3 and netcdf-4.

New Activities

Improved NetCDF/Zarr Integration

The netCDF team has now released multiple releases of netCDF-C which support the ncZarr protocol. This work has been well received, and we continue to make improvements. We are now focused on improving the S3 support for libnetcdf/ncZarr. Work continues in collaboration with the Zarr community group and the Zarr Enhancement Protocol group. **The netCDF team recognizes the need to improve messaging around the new functionality which has been implemented**, and will be working to make these features more widely known.

Over the next three months, we plan to organize or take part in the following:

- Release iterative versions of netCDF-C, netCDF-Fortran, netCDF-C++.
- Continue modernizing/editing the netCDF documentation to provide easy access to documentation for older versions of netCDF.

Over the next twelve months, we plan to organize or take part in the following:

- Release an official Windows port of the netCDF-Fortran and netCDF-C++ interfaces.
- Continue to encourage and support the use of netCDF-4's enhanced data model by third-party developers.
- Expand support for native object storage in the netCDF C library.
- Continue to represent the Unidata community in the HDF Technical Advisory Board process.
- Continue to represent the Unidata community in the Zarr/n5 collaboration conference calls.

Beyond a one-year timeframe, we plan to organize or take part in the following:

- Improve scalability to handle huge datasets and collections.
- Improve the efficiency of parallel netcdf3 and parallel netcdf4.
- Continue to add support for both file-storage and object-storage options.

Relevant Metrics

Google Metrics

Google hits reported when searching for a term such as netCDF-4 don't seem very useful over the long term, as the algorithms for quickly estimating the number of web pages containing a specified term or phrase are proprietary and seem to change frequently. However, this metric may be useful at any particular time for comparing popularity among a set of related terms.

Currently, Google hits, for comparison, are:

- **1,160,000** for netCDF-3
- **1,020,000** for netCDF-4
- **5,110** for ncZarr
- **2,230,000** for HDF5
- **174,000** for GRIB2
- **4,270,000** for ZARR

Google Scholar hits, which supposedly count appearances in peer-reviewed scholarly publications, are:

- **457** for netCDF-3
- **1,480** for netCDF-4
- **42** for ncZarr
- **45,800** for netCDF
- **26,500** for HDF5
- **1,840** for GRIB2
- **8,730** for ZARR

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Managing Geoscience Data**
by supporting the use of netCDF and related technologies for analyzing, integrating, and visualizing multidimensional geoscience data; enabling effective use of very large data sets; and accessing, managing, and sharing collections of heterogeneous data from diverse sources.
2. **Providing Useful Tools**
by developing netCDF and related software, and creating regular software releases of the C, C++ and Fortran interfaces; providing long-term support for these tools through the various avenues available to the Unidata staff (Github, eSupport, Stackoverflow, etc).

3. Supporting People

by providing expertise in implementing effective data management, conducting training workshops, responding to support questions, maintaining comprehensive documentation, maintaining example programs and files, and keeping online FAQs, best practices, and web site up to date; fostering interactions between community members; and advocating community perspectives at scientific meetings, conferences, and other venues.

Prepared April 2024

Status Report: Python

November 2023- April 2024

Ryan May, Drew Camron, Julien Chastang, Ana Espinoza, Nicole Corbin, Thomas Martin

Executive Summary

Unidata's Python efforts continue to encompass: training on the use of Python for the community; development and maintenance of several tools for the community (most notably MetPy but also Siphon and data processing scripts); and participation within the broader scientific Python community. We continue to lead and support a variety of educational efforts, including our first collaboration with the USGS on a training session. We are also furthering development of asynchronous training materials through Project Pythia, where we are working to migrate our existing workshop and gallery materials into a so-called "cookbook" within the broader project. MetPy development continues with the 1.6.0 feature release (largely calculation additions and standardization of relative humidity definition) and 1.6.1/1.6.2 bugfix releases. We are in the early planning stage for a 1.7 release which will center around min/max identification and S3 cloud data clients. Community use of the library is proving extensive, with 90 theses and peer-reviewed publications mentioning or citing MetPy in 2023; this brings the total count to over 300. We continue to assist the broader community with participation within matplotlib, cartopy, and conda-forge, though it has become increasingly difficult to dedicate time to these efforts given the full portfolio of responsibilities.

Questions for Immediate Committee Feedback

Nothing at this time.

Activities Since the Last Status Report

Python Training Efforts

Our content, expertise, and Science Gateway JupyterHub infrastructure led and supported multiple synchronous teaching/training workshops since our last update, including:

- 2024 AMS Annual Meeting Short Course on MetPy
- 2024 AMS Annual Meeting Student Conference workshop on Python
- 2-day MetPy Workshop Session Collaboration with COMET, USGS on HyTEST

We continue to support Project Pythia's Cookbook ecosystem, and the infrastructure and educational content it provides to the community. We directly engage with content development, infrastructure and maintenance, and community engagement. We are leading the organization of the 2024 Pythia Hackathon. *MetPy Cookbook* development will be a focus of the Pythia hackathon and 2024 Unidata internship.

Progress has been made on the following:

- We have iteratively implemented and updated our contained workshop content and rely on it for continued workshop offerings
- We are working with Unidata AI/ML to co-develop MetPy + ML content for Fall 2024 and AMS 2025 offerings
- We are working with Unidata Instructional Design to continue offering our Python Readiness to previously- and newly-participating institutions in Fall 2024 and are coordinating these efforts now
- We will be once more offering our summer data visualization session to interns through the UCAR Professional Development Workshop Series
- Unidata continues to be a primary collaborator on Project Pythia, particularly through support of the growing Cookbook ecosystem. Development has begun on the *MetPy Cookbook*
- John Leeman continues to lead the “MetPy Mondays” effort. We have engaged community members with plans to broaden the impact of MetPy Mondays content with code snippets and reference examples

MetPy

Development continues to be driven by requirements for our dedicated awards (in addition to bug reports and pull requests from community members). MetPy 1.6.0 was released in late December 2023, with 1.6.1 and 1.6.2 bug fix releases following up in early 2024, which included:

- Added downdraft CAPE, Galvez-Davison index, Corfidi MCS motion, and wet bulb potential temperature calculations through community contributions
- Standardized use of WMO definition of relative humidity throughout the library
- Community contributed performance enhancements to GEMPAK reading support
- Enhanced styling control in simplified plotting interface
- Support for parsing rotated pole projection metadata
- Made coordinate detection more robust to address issues with ERA5 data
- Support for dBz as parsed unit

Moving forward, 1.7 is planned for release in the Summer of 2024. Internal priorities for this release are focused on automatic identification of field maxima and minima, combined with plotting support, as well as shipping clients for S3-based datasets like NEXRAD Level $\frac{2}{3}$ and GOES. It is likely that this release will gain additional community-contributed features.

On the community participation front, the “MetPy Community Call” continues to be challenging to get off the ground. With scheduling being an on-going challenge, we are looking at trying to do such calls quarterly, which will provide higher quality discussions and presentations around the annual User Survey and AMS updates which we hope will improve interest, both internally and on the community side. Given the variety of items that are prioritized across the team’s limited time, this will likely remain challenging to prioritize overall.

Progress has been made on the following:

- MetPy 1.6.0 released late December 2023
- MetPy 1.6.1 released early January 2024
- MetPy 1.6.2 released early April 2024
- Work towards requirements of MetPy-related NSF awards
- Community awareness continues to grow, with the volume of engagement (especially support requests) and mentions on social media growing; the MetPy twitter account has reached 2831 followers (good growth despite the platform's overall instability)
- 90 theses or peer-reviewed publications cited or mentioned MetPy in 2023; this is in comparison to 62 in all of 2022. So far in 2024, there are 10 with several more awaiting full publication. We just passed 300 total theses and peer-reviewed publications that cite or mention MetPy.

Siphon and Data Processing

Siphon continues to exist in a steady state—continued maintenance and use, but minimal feature advancement. Some of this is due to limited development resources being focused on MetPy's needs; it is also due to limited pressing needs on the data access side. Largely, Siphon meets the needs we have identified for Python data access (that aren't also already met by zarr, xarray, etc.). With that said, Siphon does remain an important part of the stack used by our training work, and by Unidata's community of Python users in general. The most pressing developments we anticipate for Siphon, outside of the need for a bit of focused maintenance, are improvements to working with Siphon in interactive sessions, like the Jupyter notebook environment: improved catalog crawling interface, better string representations, and tab completion. The decision has been made to separate **non-TDS functionality** (e.g. Wyoming Upper Air archive access) out into a new remote-access toolset contained within MetPy, and we hope to begin this transition work soon.

We also continue to maintain the LDM Alchemy repository as a collection of LDM processing scripts in Python. Currently this includes the code powering the AWS NEXRAD archive as well as the program that reconstitutes NOAAPORT GOES-16/17 imagery. As we transition more of our internal data processing to Python, this repository will hold those scripts. We have seen several community questions regarding both the GOES and NEXRAD processing software.

External Participation

The Python team attends conferences as well as participates in other projects within the scientific Python ecosystem. This allows us to stay informed and to be able to advocate for our community, as well as keep our community updated on developments. As participants in a broader Open Source software ecosystem, the Python team regularly encounters issues in other projects relevant to our community's needs. As such, we routinely engage these projects to address challenges and submit fixes. We also continue to host Jeff Whittaker's netCDF4-python project repository; Jeff continues to be the active maintainer of the project. The overall involvement helps ensure that important portions of our community's Python stack remain well-supported. Ryan May continues to serve as a core developer for CartoPy as well as a member of Matplotlib's Steering Council and conda-forge's core team. *It should be*

noted, though, that it's becoming increasingly difficult to dedicate time to these efforts given the full portfolio of responsibilities on the team.

Progress has been made on the following:

- Ryan May continues involvement with the matplotlib, Cartopy, and conda-forge projects, albeit with minimal time investment.
- We also continue to actively engage with the xarray and pint projects.

Ongoing Activities

We plan to continue the following activities:

- Supporting Unidata's collection of online Python learning materials
- Engaging in synchronous Python teaching opportunities, virtual or otherwise
- Maintaining Siphon as a tool for remote data access across a variety of services
- Growing and developing MetPy as a community resource for Python in meteorology
- Continued participation in the scientific Python community as advocates for the atmospheric science community
- Working with JupyterHub as a way to facilitate data-proximate analysis
- MetPy Mondays for engaging the community

New Activities

Over the next three months, we plan to organize or take part in the following:

- Provide extensive planning and execution support for the 2024 Pythia Hackathon
- Release first version of *MetPy Cookbook* with support from community and intern development
- Plan training efforts (e.g. short course, student workshop) for 2025 AMS Annual Meeting
- Engage in continued support of Project Pythia and adjacent UCAR Python education efforts
- Attending the SciPy 2024 meeting in Tacoma, WA


Over the next twelve months, we plan to organize or take part in the following:

- Release MetPy 1.7 including high/low identification and cloud data clients
- Offer additional virtual or in-person Python workshops
- Separate non-TDS siphon capability into new MetPy remote functionality
- Explore ways to leverage Web Assembly to provide MetPy as an in-the-web-browser experience for users


Relevant Metrics

NOTE: conda-forge numbers are problematic due to an upstream known issue with "new format" packages that has hopefully been addressed for future downloads.

Siphon Downloads



MetPy Downloads



MetPy

- 96% test coverage
- Watchers: 60
- According to GitHub, 505 repositories and 57 packages depend on MetPy
- Downloads for the releases made in the last year (Conda + PyPI):
 - 1.5.0: 37397
 - 1.5.1: 110163
 - 1.6.0: 16040
 - 1.6.1: 121156
- Since 1 November 2023
 - Active Issues: 62 (46 created, 37 closed)
 - Active PRs: 148 (130 created, 139 closed)
 - External Issue Activity: 40 opened, 72 comments
 - External PR Activity: 15 opened, 22 comments
 - Unique external contributors: 41
 - Stars: 65 (1181 total)
 - Forks: 0 (377 total)
 - Commits: 357
- Since 1 April 2023
 - Active Issues: 136 (81 created, 73 closed)
 - Active PRs: 364 (344 created, 352 closed)
 - External Issue Activity: 67 opened, 152 comments
 - External PR Activity: 46 opened, 98 comments
 - Unique external contributors: 77
 - Stars: 160 (1181 total)
 - Forks: 2 (377 total)
 - Commits: 836

Siphon

- 98% test coverage
- Watchers: 27
- According to GitHub, 198 repositories and 24 packages depend on Siphon
- Since 1 November 2023
 - Active Issues: 9 (8 created, 2 closed)
 - Active PRs: 6 (4 created, 2 closed)
 - External Issue Activity: 8 opened, 7 comments
 - External PR Activity: 2 opened, 1 comments
 - Unique external contributors: 8
 - Stars: 14 (192 total)
 - Forks: 0 (65 total)
 - Commits: 0
- Since 1 April 2023
 - Active Issues: 14 (13 created, 3 closed)
 - Active PRs: 78 (68 created, 62 closed)
 - External Issue Activity: 11 opened, 15 comments

- External PR Activity: 4 opened, 6 comments
- Unique external contributors: 12
- Stars: 27 (206 total)
- Forks: 0 (67 total)
- Commits: 17

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Providing Useful Tools**

Python has become a key tool in the atmospheric sciences, and the geosciences in general. MetPy leverages the rest of the scientific Python ecosystem to provide a suite of documented and tested domain-specific functionality, supporting greater use of Python by the community. Siphon serves to provide access to the growing collection of remote data sets. Together, MetPy and Siphon give the community a platform for scripted analysis of real-time and archived weather data. These tools are also readily used in the Jupyter Lab/Notebook environment, for ease of use in cloud and HPC computing environments, facilitating data-proximate analysis. We also participate in a variety of projects in the broader scientific Python ecosystem, to help ensure the ecosystem's viability and that it continues to meet our community's needs.

2. **Supporting People**

We provide a variety of online training resources to facilitate our community's education and use of Python. We also regularly conduct training workshops to teach attendees how to use tools and apply them to their problems and challenges in research and education.

Prepared April 2024

Status Report: Support

November 2023- April 2024

Jennifer Oxelson, UPC Staff

Executive Summary

Unidata staff have started vetting replacement packages for our current eSupport system. We hope to have a new support package in use by the end of Spring 2024.

The AWIPS and THREDDS developers have performed a staggering amount of support for their packages this past year (820 and 549 support replies respectively). 🙌

Questions for Immediate Committee Feedback

What changes/recommendations would you like to see in Unidata support or how Unidata conducts support? Are you subscribed to any of the unidata software package mailing lists?

Activities Since the Last Status Report

Retirement of Tom Yoksas

Tom Yoksas has left for greener pastures (retirement), but will still be volunteering with Unidata Support efforts when applicable.

McIDAS development at Unidata has been curtailed, but support for the current software will continue for the immediate future. Unidata will continue to license the software on behalf of the community (via a MUG membership) for the foreseeable future as well.

Looking for a new in-house support package

- The currently-used eSupport package is long in the tooth. Unidata staff have started vetting replacement packages. We hope to have a new support package in use by the end of Spring 2024.

Training

Unidata training/workshop information can be found in the [Community status report](#).

New Activities

In order to fulfill our objectives articulated in the Unidata 2018 Proposal, focused efforts are needed in two major areas:

- Enhance electronic support offerings
- [Create instructional materials for online virtual training](#)

Relevant Metrics

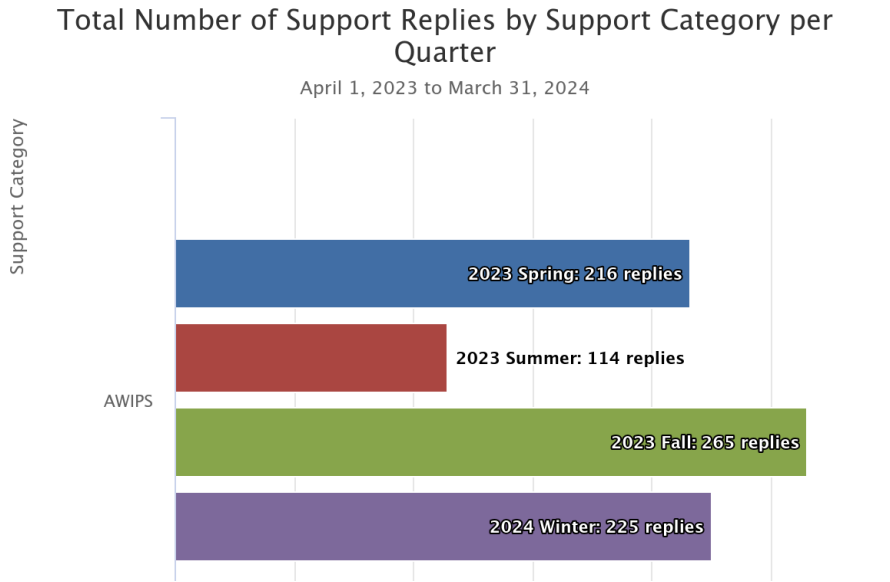
User support "transactions" (new inquiries and follow-ups) have been processed through the Unidata inquiry tracking system (eSupport). Other methods of providing answers to questions posed (e.g., Github, Stack Overflow, Jira, mailing list replies, etc.) add substantially to the support load.

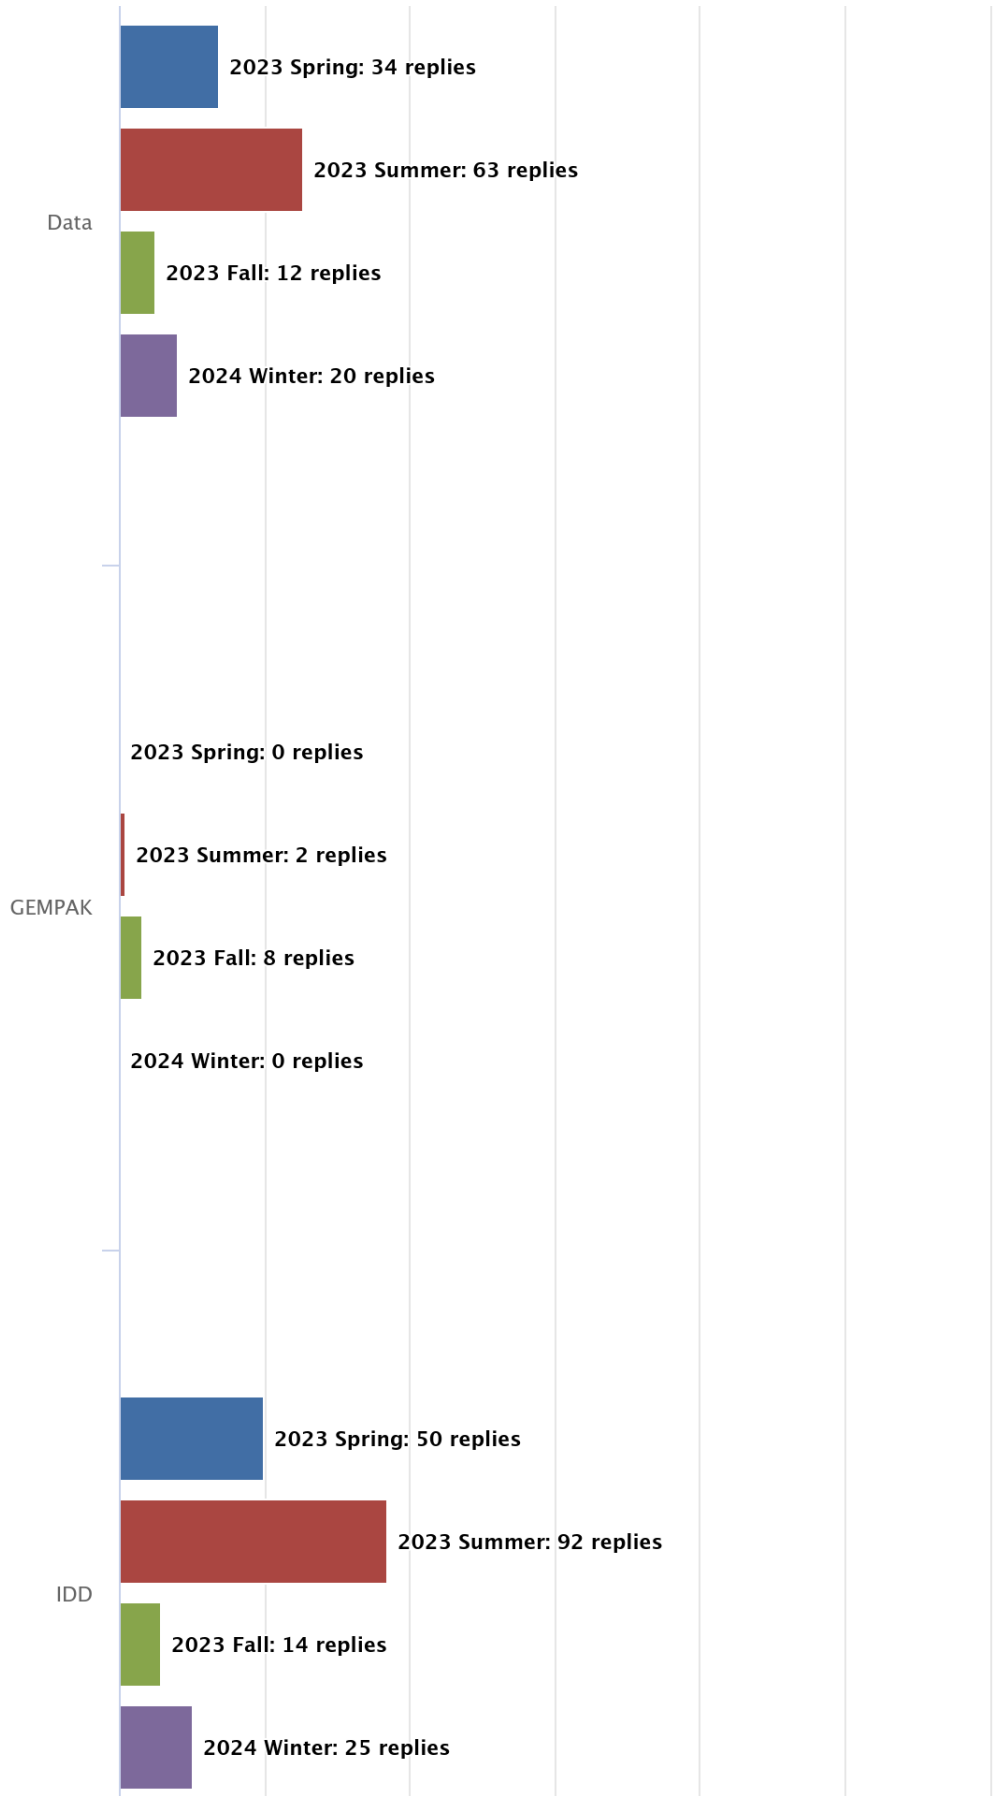
Additional metrics may be found in the [Comprehensive Metrics Data](#) portion of this meeting's agenda.

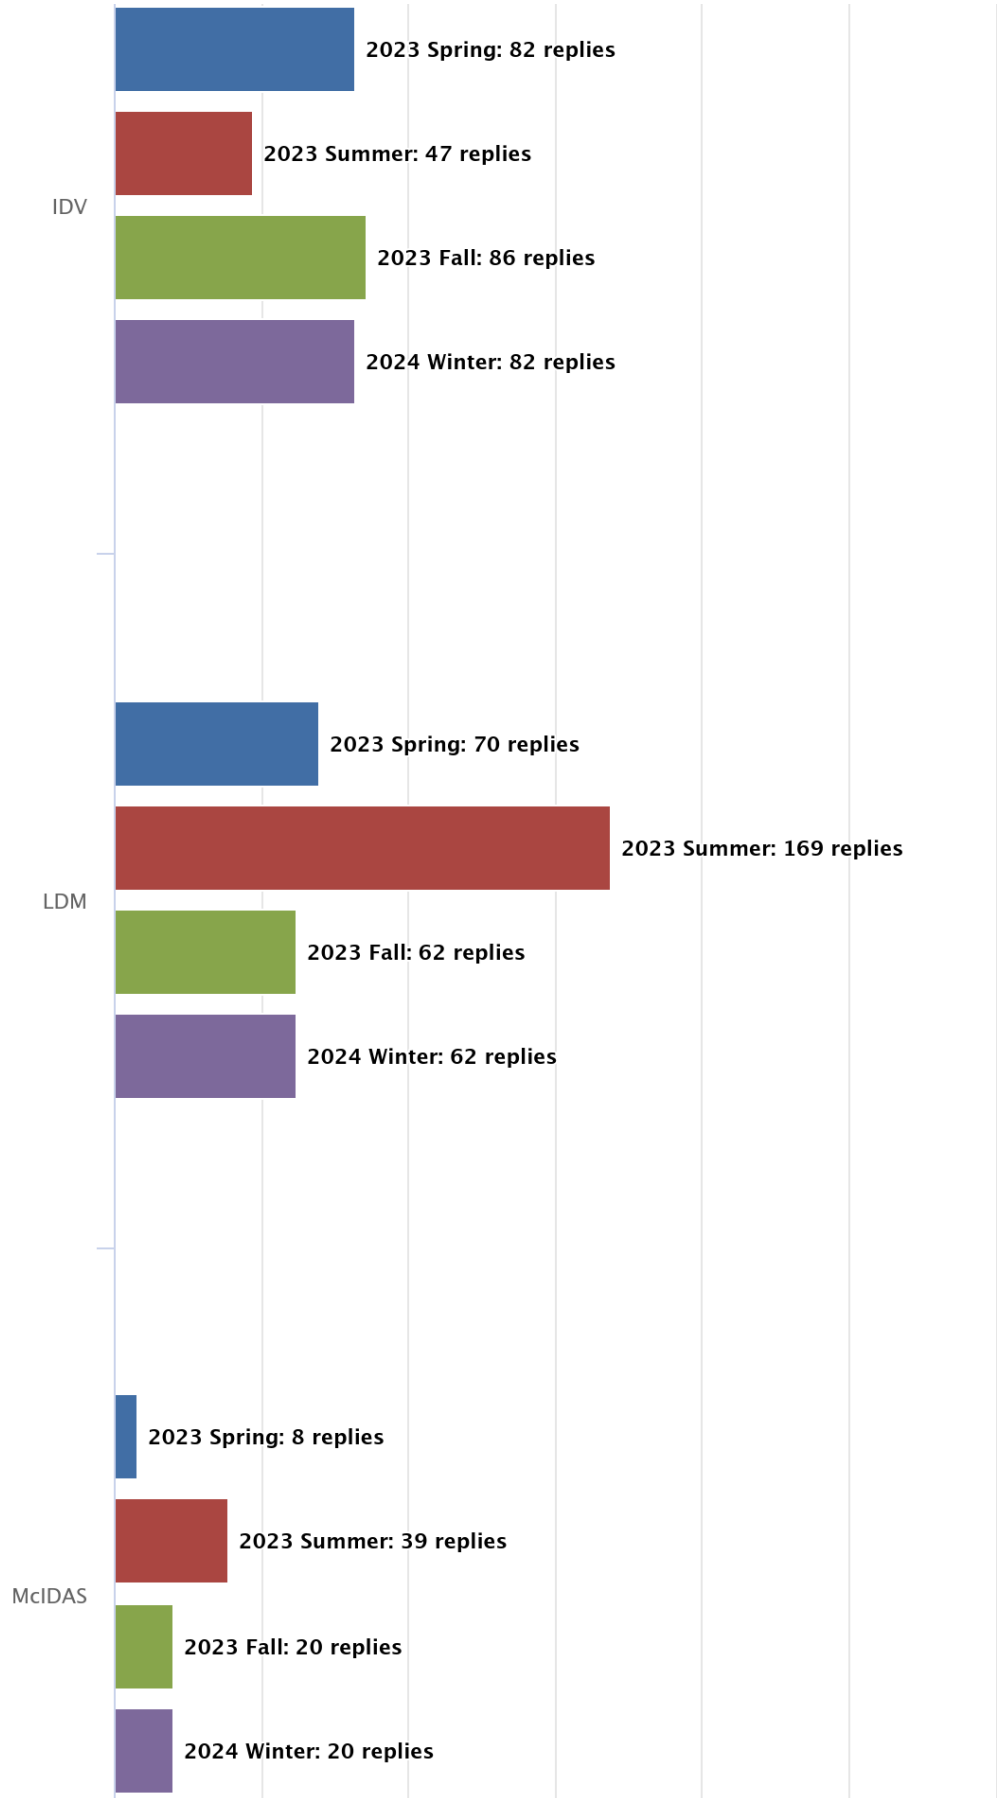
Fig. 1: Below are histograms that portray the number of Unidata email responses for categories of support logged in the Unidata Inquiry Tracking System for the 12 month period from **October 1, 2022 until September 30, 2023**.


The quarters shown are defined as:

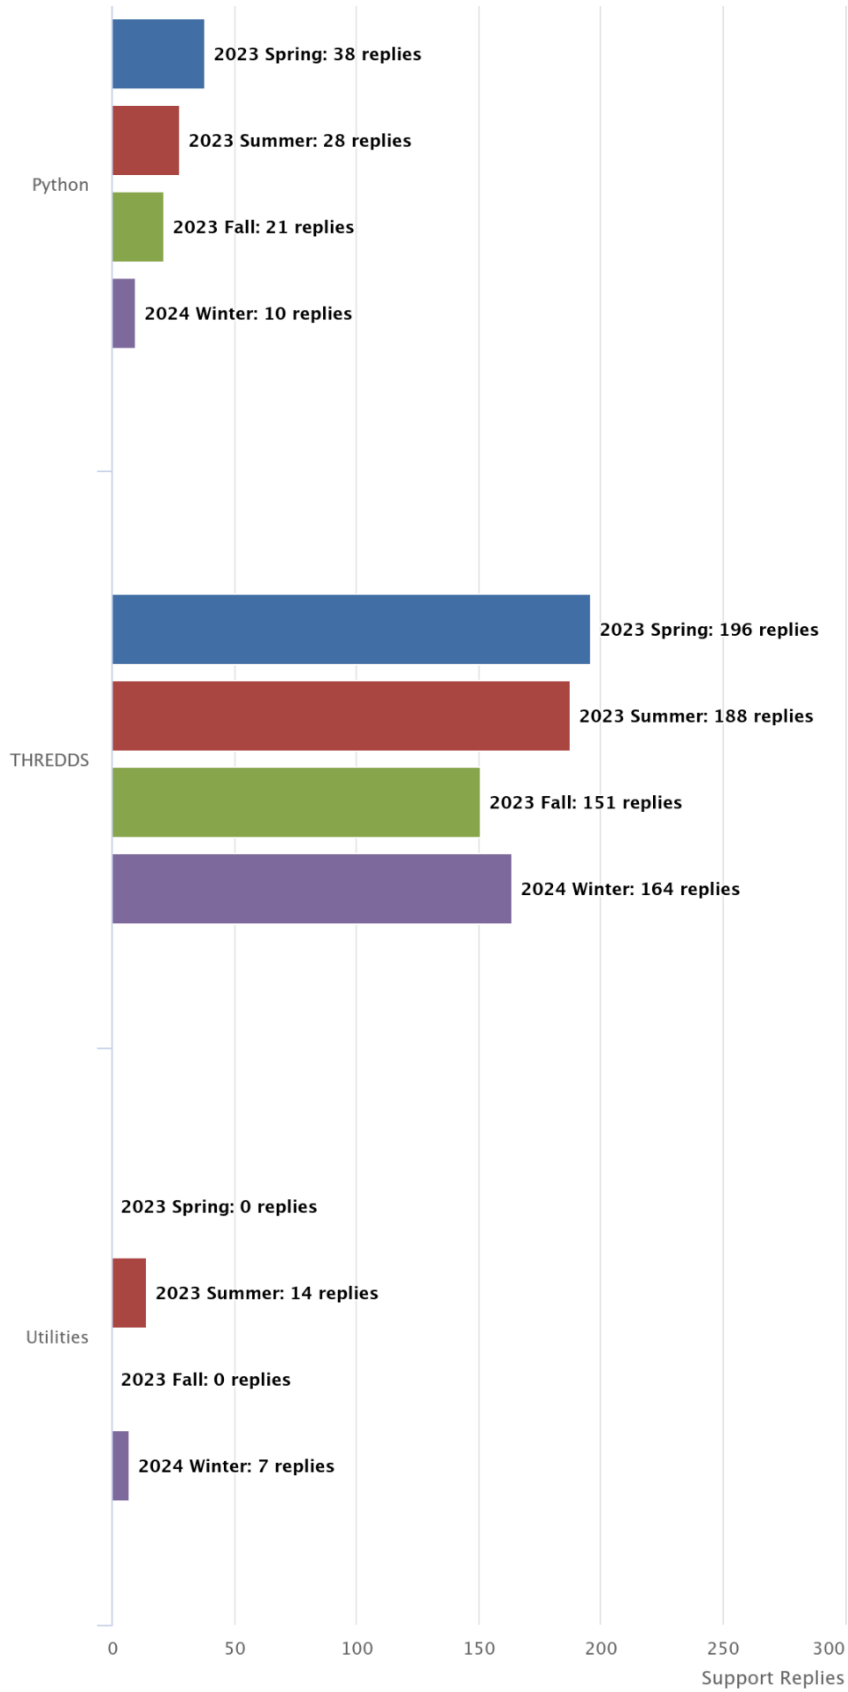
Winter: January, February, March **Spring:** April, May, June **Summer:** July, August, September **Fall:** October, November, December











Click an item (below) to hide it's data from the chart above

- 2023 Spring
- 2023 Summer
- 2023 Fall
- 2024 Winter

Individual support activities included in the categories shown above are listed in the following table.

Category	eSupport Departments
AWIPS	Support AWIPS
Data	Support CaseStudy, Support CONDUIT, Support Datastream, Support LEAD, Support Level II, Support NOAAPORT, Support SUOMINET
GEMPAK	Support GEMPAK
IDD	Support IDD, Support IDD Antarctica, Support IDD Brasil, Support IDD Cluster, Support IDD SCOOP, Support IDD TIGGE
IDV	Support IDV, Support IDV Storm, Support McV, Support VisAD
LDM	Support LDM
McIDAS	Support McDevelop, Support McIDAS, ldm-decoders
Miscellaneous	Administration, Development, Plaza, Staging Folder, Support, Support eSupport, Support Miscellaneous, Support Platforms, Support Plaza, Student Interns, Systems, Support Machine Learning
NetCDF	Support LibCF, Support netCDF
Outreach	Outreach, Polcomm, Science Gateway, Support Egrants, Support eLearning, Support News, Support Outreach, Support Workshop, Usercomm, Student Interns
Python	Support Python
RAMADDA	Support RAMADDA
THREDDS	Support netCDF Java, Support THREDDS
Utilities	Support LDM-McIDAS, Support netCDF Decoders, Support netCDF Perl, Support OPeNDAP, Support Rosetta, Support UDUNITS

Comments

- The total support provided by the UPC continues to be substantial: yearly totals have shown a slight decline over the past several years, but this is most likely attributable to the increased ways support is being provided. Overall support activities vary by somewhat by quarter. Spikes in support for individual packages is largely correlated with the releases of new distributions of the packages.
- GEMPAK is now fully community supported, with occasional input from Unidata staff.
- McIDAS development at Unidata has been curtailed, and support for the current software will continue for the immediate future. Unidata will continue to license the software on behalf of the community (via a MUG membership) for the foreseeable future as well.
- The AWIPS and THREDDS groups have answered a staggering amount of support requests this past year (820 and 549 support replies respectively).
- The Winter 2024 quarter has seen a big increase in support questions answered in the Miscellaneous category, resulting from the systems staff utilizing the previously underused support-systems department. These support requests are mainly from Unidata staff, but also include requests from community members.

Notes

These numbers and conclusions should not be taken too literally, for several reasons:

- For some packages, multiple responses in the same thread may be bundled into a single archived email. Other packages have each response in a thread counted separately.
- After a new release of software, there may be a flurry of the same or similar questions, which can be answered in separate emails or in a single mailing list posting.
- The graph primarily represents support of end users and site administrators, not developers. Support for non-Unidata developers in projects such as THREDDS, IDV, etc., requires significant resources, but is difficult to assess.
- Not all support records were indexable for this report. Given this, the above numbers are an ****underestimate**** of the actual support being provided by the UPC.

[Additional User Support Metrics](#)

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Managing Geoscience Data**

Unidata User Support enables access to geoscience data by supporting the use of tools created and/or supported by the UPC.

2. **Providing Useful Tools**

A significant part of providing useful tools is providing support for those tools. Unidata has always provided world class support for all of the tools that it makes freely available to the greater geoscience community.

3. **Supporting People**

The user support provided by the UUPC is recognized throughout the atmospheric science community. Unidata's outreach efforts are routinely noted as being exceptional in surveys of the NCAR/UCAR community.

Prepared *April 2024*

Status Report: THREDDS

November 2023- April 2024

Hailey Johnson, Tara Drwenski, Jennifer Oxelson, Ryan May, Ethan Davis

Executive Summary

The deprecation of the Spring 5.x library sets a hard deadline for a significant number of upgrades in the TDS by December 2024. We plan to, at that time, release a long-term maintenance version of the TDS, before shifting focus to a replacement data service system.

Questions for Immediate Committee Feedback

None at this time.

Activities Since the Last Status Report

General

- Megan Lerman left the THREDDS team due to funding considerations.

NetCDF-Java

- Improvements to performance and accuracy in the application of dataset enhancements.
- Removed Gretty as a dependency from the project; implemented tests using Test Containers instead.
- Wrapping up the release of version 5.6.0.

TDS

- Version 5.5 will be released soon as a feature-complete release with improved performance for WMS and NCSS.
- The performance testing pipeline is live, and performance trends can be observed real time on the Unidata status page.
- The TDS has been separated from the THREDDS ISO services, which reduces the overhead involved in the THREDDS release process. ISO services must now be explicitly included by the user.

Ongoing Activities

Server management

- We have continue to maintain our three primary Unidata THREDDS servers:
 - thredds.ucar.edu always runs the latest stable release of the TDS (unless a quick security update is required)
 - thredds-test.ucar.edu automatically deploys new versions of the TDS when netCDF-Java or the TDS GitHub repositories update; it is therefore always running the latest development version
 - thredds-dev.ucar.edu is intended to be used by THREDDS developers, rather than THREDDS users; we use this domain to test changes that require access to “real” data
- The thredds-rc and thredds-jumbo hostnames have been retired.
- Additionally, the Cloud Activities group manages cloud-hosted TDS instance (see cloud activities report for details).

Maintenance

- Maintain thredds.ucar.edu and continue to add new datasets as they become available or are requested by users.
- Closely monitor the security status of our project dependencies, and provide updated versions of our libraries and server technologies to address as needed.
- Continue to respond to user feedback regarding TDS 5.x and transitioning servers to the latest version.
- Communicate with the TDS users about the upcoming infrastructural updates and what will need to be done by TDS administrators.

Development

- With the exception of a last few bug fixes in the TDS and infrastructure upgrades, development in the THREDDS libraries is on hold while we consider future directions.

New Activities

Over the next three months, we plan to organize or take part in the following:

- netCDF-Java
 - Release version 5.6.0
- TDS
 - Release version 5.5 of the TDS.

Over the next twelve months, we plan to organize or take part in the following:

- netCDF-Java
 - Upgrade the version of Gradle used in the build
- TDS
 - Upgrade the following libraries:

- Gretty 3 - Gretty 4
- Tomcat 8 and 9 - Tomcat 10
- Gradle 6 - Gradle 7
- Java 11 - Java 17
- Spring 5 - 6

Beyond a one-year timeframe, we plan to organize or take part in the following:

- netCDF-Java
 - Determine the future of the library; should there be a version 6?
- TDS
 - Maintain a secure THREDDS server at thredds.ucar.edu
 - Develop a replacement for Unidata to provide ESS data to the community.

The following active proposals directly involve THREDDS work:

- Partnering with the netCDF team on a proposal to fund the development of shared, modernized build and CI infrastructure across the netCDF libraries.

Relevant Metrics

THREDDS Startup Metrics


	2023-03 — 2023-10	2014-08 — 2023-10
TDS Startup (unique IP address count)	2185	44219
	Total Servers	Information page updated
Publicly Accessible ¹ TDS count	119	68


Over the past 4 months, ****2,185**** unique IPs started up the TDS (November 2023 through April 2024). Since we’ve started tracking these metrics (v4.5.3, August 26th, 2014), we’ve seen the TDS startup from ****44219**** unique IP addresses. There are currently ****135**** publically accessible TDSs running “in the wild”. Of the ****119**** publically accessible servers, ****68**** have updated the name of their server in their server configuration file (taken as a sign that they are maybe, possibly, intended to be used by others...maybe...).

The figures below show the distribution of TDS versions (top), and the fractional share of servers running version X or older (bottom). Each labeled version includes betas and snapshots, not just the official release of that version, for presentation simplicity. TDS v5.0

¹ “Publicly accessible” means we could find a top-level THREDDS Client Catalog. We checked <server>/thredds/catalog.xml (version 4), <server>/thredds/catalog/catalog.xml (version 5), including the most common ports of 80, 8080, 443, and 8443.

remains the dominant specific version running in the wild, although there are almost as many instances of 5.4.





Strategic Focus Areas

The THREDDS projects covered in this report support the following goals described in Unidata Strategic Plan:

1. Managing Geoscience Data

The component software projects of the THREDDS project work to facilitate the management of geoscience data from four points of view: __Making Geoscience Data Accessible, Making Geoscience Data Discoverable, Making Geoscience Data Usable, and Enhancing Community Access to Data__. As a client-side library, **netCDF-Java** enables end users to read a variety of data formats both locally and across numerous remote technologies. Less user-friendly formats, such as GRIB, are augmented with metadata from community driven metadata standards (e.g. Climate and Forecast metadata standards), and viewed through the more user friendly Common Data Model (very similar to the netCDF Data Model), providing a single set of Java APIs for interacting with a multitude of formats and standards. The **THREDDS Data Server** exposes the power of the netCDF-java library outside of the Java ecosystem with the addition of remote data services, such as __OPeNDAP__, __cdmremote__, __OGC WCS__ and __WMS__, __HTTP direct download__, and other remote data access and subsetting protocols. The TDS also exposes metadata in standard ways (e.g. ISO 19115 metadata records, json-ld metadata following schema.org), which are used to drive search technologies. **Rosetta** facilitates the process of translating ascii based

observational data into standards compliant, archive ready files. These files are easily read into netCDF-Java and can be served to a broader community using the TDS.

2. **Providing Useful Tools**

Through Rosetta, the THREDDS project seeks to intercede in the in-situ based observational data management lifecycle as soon as possible. This is done by enabling those who produce the data to create archive ready datasets as soon as data are collected from a sensor or platform without the need to write code or intimately understand metadata standards. NetCDF-java and the TDS continue to support legacy workflows by maintaining support for legacy data formats and decades old data access services, while promoting 21st century scientific workflows through the creation of new capabilities and modernization of existing services (e.g. Immutability, upgraded technical stack, microservice development).

3. **Supporting People**

Outside of writing code, the THREDDS project seeks to support the community by __providing technical support, working to build capacity through Open Source Software development, and by building community cyber-literacy__. The team provides expert assistance on software, data, and technical issues through numerous avenues, including participation in community mailing lists, providing developer guidance on our GitHub repositories, and leading and participating in workshops across the community. The team also actively participates in “upstream” open source projects in an effort to help sustain the efforts of which we rely and build upon. We have mentored students as part of the Unidata Summer Internship Program, and worked across organizations and disciplines in support of their internship efforts.