Spring 2026 Users Committee Meeting

Virtual Meetings (All times are Mountain Standard Time)

Monday, 16 March 2026, 9:00 - 11:00 MT

09:00 – 09:15 Administrative Items and Welcome/Around-the-table check-ins (Co-Chairs/Doug Dirks)

09:15 – 09:45 Director's Report (Mohan Ramamurthy)

09:45 – 10:15 NOAA NCEP Update (Margaret Curtis)

10:15 – 10:45 IDV and Gemini Demo and Q&A (Yuan Ho)

10:45 – 11:00 Administrative Items / Wrap Up

Wednesday, 18 March 2026, 10:30 - 12:30 MT

10:30 – 10:35 Administrative Items and Welcome (Co-Chairs/Doug Dirks)

10:35 – 11:15 NSF Unidata Portfolio Lightning Updates (UPC)

11:15 – 11:30 BREAK

11:30 – 12:15 NSF Unidata Portfolio Open Discussion (all)

12:15 – 12:30 Administrative Items / Wrap Up

Friday, 20 March 2026, 13:00 - 15:00 MT

13:00 – 13:05 Administrative Items and Welcome (Co-Chairs/Doug Dirks)

13:05 – 13:20 DeSouza Candidate Selection and Discussion

13:20 – 13:50 Invited Speaker (TBD)

13:50 – 14:05 Future Directions for Earth Observations & Data Stewardship (Sean Arms)

14:05 – 14:45 Open Discussion: Shifting Landscapes at Higher Education Institutions

14:45 – 15:00 Administrative Items / Wrap Up

*Friday, 10 April 2026, 12:30 - 15:30 MT
(Joint) + Optional Mingling

12:30 – 12:45 Administrative Items and Welcome (Co-Chairs/Tanya Vance)

12:45 – 13:45 Committee Model Implementation and Chair Selection

13:45 – 14:00 BREAK

14:00 – 14:45 AI/ML Vision and Opportunities for NSF Unidata (Subcommittee, UPC)

14:45 – 15:15 WIS2 Overview and Discussion (Ethan Davis)

15:15 – 15:30 Administrative Items / Wrap Up

15:30 – 16:30 Virtual Mingling (all, optional)

Status Report: Users Committee Actions

September 2025 - March 2026

Unidata Program Center Staff

Actions from the Previous Meeting (September 2025)

Action 1

Define AI/ML from the community perspective. Determine where Unidata's "lane" is in the AI/ML space. Committee to have something to share with UPC by mid-November. [SeanF, KimW, AaronK, CraigR]

Result

Complete

Action 2

Poll committee members for Spring 2026 meeting dates (virtual meeting) [TanyaV]

Result

Complete

Action 3

Match Status Report order with Lightning Talks and add Q&A slide to end of Lightning Talks that includes any Status Report questions [TanyaV, DougD]

Result

This change is implemented in the Executive Summary and in the full status report PDF. The ordering of the reports on the meeting web site will remain unchanged.

Action 4

Add an agenda item for spring meeting for IDV & Gemini demo [TanyaV, Yuan]

Result

Complete

Action 5

Inquire about virtual townhall options for students at AMS [ScottS]

Result

Evaluate if desired AMS 2027

Prepared *March 2026*

Status Report: Community Services

September 2025 - March 2026

Doug Dirks ,Tanya Vance

Executive Summary

The Community Services group has predominantly focused on progressing downstream phased activities for the prioritized Reimagined Science Gateway Reimagined/Education Hub Project, Community Assessment Initiative, and NSF Unidata Website, as well as coordinating and participating in community engagement and outreach, supporting program development proposals and activities, and advancing the COMET Learning Management System (LMS) MetEd Initiative.

During this cycle, Community Services also drafted and published communications, facilitated the 2026 NSF Unidata Community Equipment Awards, and supported NSF Unidata Advisory Committee activities and communications, including follow-up actions and reports from the Fall Joint Committee meeting workshop sessions, progressing Spring meeting planning, the 2026 NSF Unidata Community Equipment Awards, committee model implementation, the final Committee Charter, and solicitation of the Russell L. DeSouza Award, as well as supported cross-program and cross-organizational collaborations and activities.

We are also excited to announce that we honored Community Services member Doug Dirks' 15 Years of Service at the UCAR Celebration of Excellence.

Activities Since the Last Status Report

Community Outreach and Services

Community Communications:

- Posts to the News@Unidata blog appear regularly, but not on a specific schedule. Some highlights:
    - IDV 7.0u1
    - Brian Blaylock Receives 2025 DeSouza Award
    - 2025 Community Equipment Awards
    - 2025 NSF Unidata Community Assessment Survey: Key Findings
    - NSF Unidata Staff at AGU Fall 2025 Meeting
    - Call for Proposals: NSF Unidata 2026 Community Equipment Awards
    - FAIR in ML, AI Readiness, & Reproducibility (FARR) Workshop
    - NSF Unidata 2026 Summer Student Internships Available!
    - NSF Unidata Staff at AMS 2026 Meeting
    - 2026 DeSouza Award Nominations

    - Unidata Artifacts Server Changes
    - AI-driven Global Model Output Availability
    - University Allocations of NSF NCAR HPC Resources Available
    - 2026 EdEC ASP Summer Colloquium
    - NSF Unidata Science Gateway Update: What's new? What's next?
- Community meetings and other announcements
- Updates to NSF Unidata's social media channels (Facebook, Bluesky, LinkedIn) and UCAR communication channels
- Publishing short videos on the Unidata YouTube channel
- Writing support for conference papers, student intern materials, and proposals (see Program collaborations and activities)

NSF Unidata Web Development and Maintenance

The Program Center launched its new website in early September, 2025. While the launch was a major milestone in what has been a long process of transitioning Unidata's public-facing web assets from systems designed and controlled by Program Center staff to a system designed and controlled by UCAR-level IT staff, getting the site up and running was not the end of the story. Firstly, only a subset of the historical materials hosted on our legacy website were in place on the new site at launch time. We continue to transfer useful historical materials to the new site; we face challenges here both from the bespoke technologies used in the legacy site and the sheer volume of material to evaluate and transfer.

Secondly, over the course of several months following the launch, we experienced significant traffic spikes (one to two orders of magnitude higher than normal), with connection levels so high that the (off-site, hosted) servers could not keep up. Our analysis of the traffic suggests that the bulk of it came from systems attempting to crawl and display pages from the legacy site; the fact that many resources have different locations on the new site (or are not yet present there) may have increased the load on the system. We also see strong suggestions that our historical support archives have been the target of Large Language Model training. While it isn't clear that the systems hitting our site were doing so with malicious intent, the effect was similar to a Distributed Denial of Service (DDoS) attack.

It took us (Unidata staff, UCAR IT staff, and technicians at the hosting provider) several attempts to settle on a series of site changes to remediate the high-traffic problem. While we believe the situation is now under control, we do still see traffic spikes occasionally.

Outreach to the Earth Systems Science community at conferences, workshops, events, and working groups

- Revamped NSF Unidata swag and created program outreach slide loop (DD)
- Coordinated and participated in engagement at UCAR Members Meeting Reception and Common Interest Tables (TV, October)
- Presented a career talk and engaged with senior undergraduates at Metropolitan State University of Denver as an invited guest speaker (TV, November)
- Coordinated AGU outreach at AGU UCAR Community Programs (UCP) Booth (TV, December)

- Coordinated and participated in outreach at AMS UCAR UCP Science Showcase Booth (TV, January)
- Gave an oral presentation, *Voices of the Earth Systems Science Community: Insights from the 2025 NSF Unidata Assessment Survey*, during the AMS EIPT Session - FAIR and Open Data and Software within the Atmospheric and Ocean Sciences to Support Transparent, Reusable, and Efficient Research and Operations II (TV, January)

NSF Unidata Community Awards

- Finalized and announced the 2025 Community Equipment Awards Selection Panel and process

2025 NSF Unidata Community Equipment Award Recipients

- Facilitated the 2026 Community Equipment Awards process, including implementing revisions recommended by the 2025 subcommittee

Call for Proposals: 2026 NSF Unidata Community Equipment Awards

DeSouza Award

- 2026 DeSouza Award Nominations

Community Engagement Initiatives

Community Assessment Initiative

- Presented to committee members and facilitated MoSCoW (Must, Could, Should, Would) prioritization planning workshop session and business development workshop session at Joint Fall Committee Meeting
- Created Key Findings Report and published key outcomes to the broader community
- Presented an oral talk at AMS on the findings and NSF Unidata actions
- Using the insights from the 2025 Community Assessment Survey as a data point for program planning and actions

Program collaborations and activities

- Coordinated NSF Unidata learning and engagement event tracking and reporting
- Supported development and drafting of NSF CSSI proposals for UDUnits and Dashboards and internal Presidential Strategic Fund WIS2 proposal
- Supported participation resources for NASEM Future Directions for Earth Observations & Data Stewardship workshop
- Progressing engagement and strategic program outreach and business development opportunities and mechanisms with UCP Strategic Partnerships Lead
- Student Summer Internship Selection Panel participant

Supported Users and Strategic Advisory Committee activities and communications

- Facilitated Spring Committee Meeting planning with members, representatives, and NSF Unidata staff
- Facilitated communications and documentation for program awards and Committee Meeting actions, including AI/ML subcommittee position, value proposition for NSF Unidata to UCP Director for outreach, summary reports from the Joint Fall Meeting workshop sessions, progressing brainstorming ideas with UCP Strategic Partnerships Lead

Progressed NSF Unidata's NSF Reporting

- Preparation of our Annual Report for NSF award #2403649 NSF (the Unidata core award) is underway. We expect to submit the report to NSF in early April, 2026.

Cross-program and cross-organization collaborations and activities

- Support UCP Collaborative for Experiential Learning (TV)
- Ambassador on NSF NCAR Convergence Science Program (CSP) Community of Practice (CoP) (TV)
    - Member Participant of NSF NCAR CSP Data Outputs group
    - Member Participant of NSF NCAR CSP Interview Analysis group
    - Selection Review Panelist for 2026 NSF NCAR EdEC Advanced Study Program (ASP) Summer Colloquium: Beyond Boundaries: Training Earth System Scientists and Engineers in Convergence Research for Compound Hazard Resilience
- Active Member on UCAR Community Engagement Group (TV)
- Active Member on UCAR Exhibit Working Group (TV)
- Active Member on UCAR Tutorials & Workshops CoP (TV)
- Interviewee in UCP Marketing Initiative (TV)
- Mentor for UCAR Mentorship Program (TV)
- Active Member on UCAR Web Technical Governance Advisory Group (WeTAG) (DD)

Learning Services

Reimagined Science Gateway / Educational Hub

- Launched Science Gateway with Integrated Educational Hub (See Cloud Computing Activities report for more details)
- Supported Data Proximate Computing and News pages and downstream phase planning
- Collaborated with COMET on Initiative to have NSF Unidata eLearning hosted on their learning management system (LMS), MetEd (in progress)

Learning offerings

- See AWIPS and IDV Activities report for details on IDV presentation and remote training demo session for AWIPS CAVE and python-awips using JupyterHub at Millersville Satellite Meteorology Class
- See AWIPS and MetPy Activities report for details on AMS Student Workshop 2026 on

Exploring Meteorological Datasets with NSF Unidata AWIPS and MetPy

Ongoing Activities

We plan to continue the following activities:

- Ongoing activities related to above (community communications, support for advisory committee activities, facilitating awards, outreach and engagement, supporting cross-program, cross-organization, and external collaborations, learning services, etc.)
- Participation in Working Groups noted above
- Support the pursuit of funding and bringing greater public awareness to NSF Unidata

New Activities

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

- Facilitate 2026 NSF Unidata Community Equipment Awards selection and process
- Facilitate spring Joint and Strategic Advisory Committee meetings and actions
- Updating eLearning modules and supporting transfer and communications of NSF Unidata eLearning resources to be hosted on COMET MetEd
- Continue to enhance NSF Unidata's new website and discoverability, including content migration, workflow-centric organization, and maintenance

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

- Progress implementation of restructured advisory committee model
- Continued maintenance and downstream phased activities for the Reimagined NSF Unidata Science Gateway Project, including on-demand notebooks and community contributions
- Proposals and program development activities
- Cross-program and cross-organizational collaborative activities
- Work to broaden our awareness, applications, reach, and impact across Earth System Science communities

Relevant Metrics

In the period following the launch of NSF Unidata's new website, we saw a sustained period of very high (an order of magnitude higher than usual) traffic, most likely from crawlers and other non-human actors. (See NSF Unidata Web Development and Maintenance, above, for more.) As a result, our website traffic metrics for the six-month period between September 2025 and February 2026 are not indicative of actual site value. We believe we should be able to resume reporting useful metrics for the next six-month period.

Social media statistics, February 26, 2025

1. # of Bluesky followers: 486 (up from 459)
2. # of Twitter/X followers: 2003 (down from 2020) (not actively posting to X)
3. # of Facebook followers: 992 (up from 982)
4. # of YouTube subscribers: 4470 (up from 4380)
5. # of LinkedIn followers: 567 (up from 492)

Unidata eLearning statistics

NSF Unidata is in progress of transferring eLearning content to COMET's MedEd LMS. This new platform will allow for enhanced metric tracking for our educational resources and learners. We believe we should be able to resume reporting metrics for the next six-month period.

Prepared *March 2026*

Status Report: AWIPS

September 2025 - March 2026

Tiffany Meyer

Executive Summary

We currently have one build (version 23.4.1) available to support RHEL/Rocky 8 and 9 that has been in production since July 2025. We have decommissioned all older builds due to running on old RHEL7 systems. EDEX, CAVE, and python-awips are available for install as well as source code available. Next release I am working on will focus on updating Java from version 11 to 17.

The National Weather Service (NWS) has started their transition to Cloud AWIPS by awarding two new 5-year contracts. Booz Allen Hamilton has been contracted for the “Data Environment” path which will focus on how AWIPS will receive, organize, store, and analyze data in a cloud based environment. Unidata will be working with BAH to complete specific Task Orders. Accenture Federal Services was awarded the “Applications Environment” contract which focuses on designing/developing data intensive user interface that serves as 1) interactive data analysis and 2) automated creation for public facing products in a cloud based environment. The goal for the NWS is to have a fully operational Cloud AWIPS environment by late 2027. Minimal new functionality is being added during this time to ensure a smooth transition, however Hazard Services V4, which includes convective hazards, is being pushed out very soon and will eventually replace WarnGen.

A couple workshops/demos were done over the past few months. I provided a remote demo of CAVE, focusing on satellite data for Millersville University. At AMS 2026 in Houston, TX, Drew and I hosted a Student Workshop focusing on data visualization within metPy, python-awips, and CAVE. We had over 60 students attend. Both of these events utilized the Virtual Desktop on a JupyterHub where students could connect to CAVE from their browser via the help of the Science Gateway group and Jetstream2.

Questions for Immediate Committee Feedback

Are there any new datasets, visualizations, capabilities you would like to see added to Unidata’s AWIPS?

Activities Since the Last Status Report

AWIPS

Our EDEX servers have been continuously running on the new Jetstream2 platform. Since the last status report, no major updates have been released. All python-awips notebooks have

been continuously updated to work with the latest packages. The main focus has been on adding new products as they come available and add new functionality on the back-end to better track our users.

A couple new products have been added to AWIPS:

- Canadian Radars (individual and mosaic)

Our blog series, AWIPS Tips, successfully ran every other week for just over three years through 2024. The breakdown of all the entries can be found on our documentation website in the Educational Resources page. Announcements and important information is shared through our mailing list (awips2-users@unidata.ucar.edu), and our social media accounts (Facebook, Bluesky, LinkedIn, and YouTube when applicable). However, due to reduced staffing levels blogs are only released on certain occasions.

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.

The AWIPS team has started tracking the number of users that are connected to our public EDEX. These numbers are likely an underestimate because any user who uses our Virtual Machine option comes across as “awips@localhost” which is only counted once. In addition to tracking the number of users, all new installs of EDEX will ping Unidata’s servers when the process starts up to gain a better understanding of who is running their own EDEX servers. We are hoping to map IP’s to Universities, private or government entities, or the general public. Below is a graph illustrating the number of daily users that have connected to our public EDEX from 20250101 - 20251231.

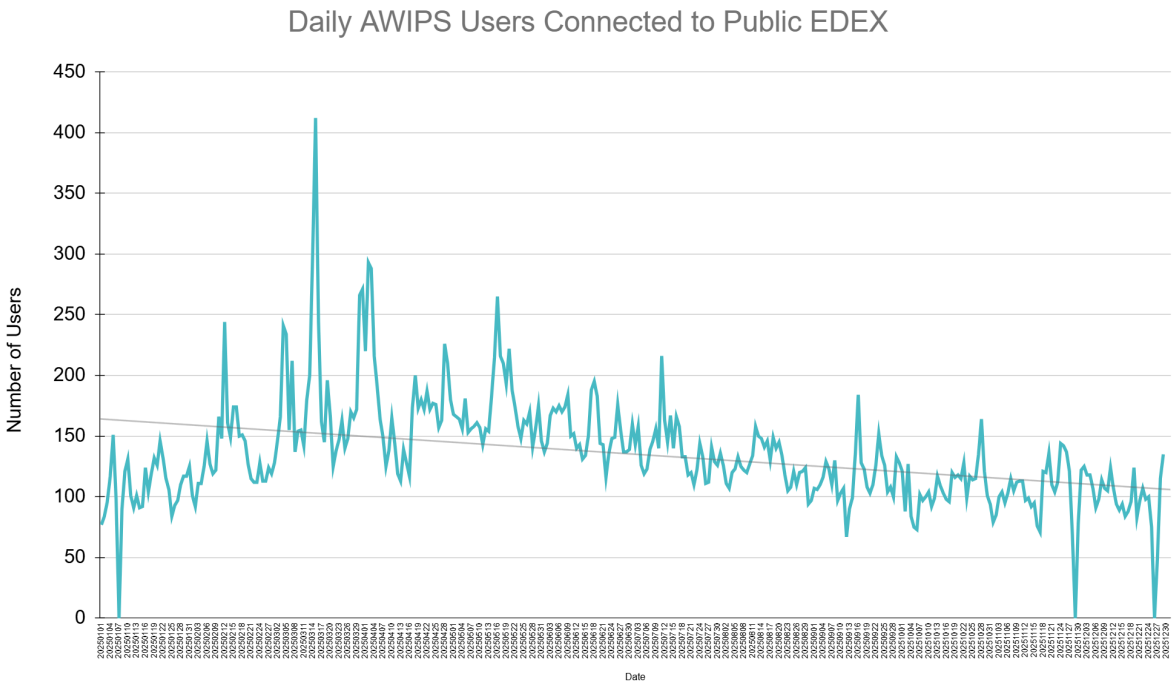

The AWIPS group typically has quite a bit of interactions with users via our support email. Through our support signatures we still have an active [support evaluation survey]. The majority of our feedback has been overwhelmingly positive, and the graphic below is a summary from all responses we've received regarding the quality of service we provide:

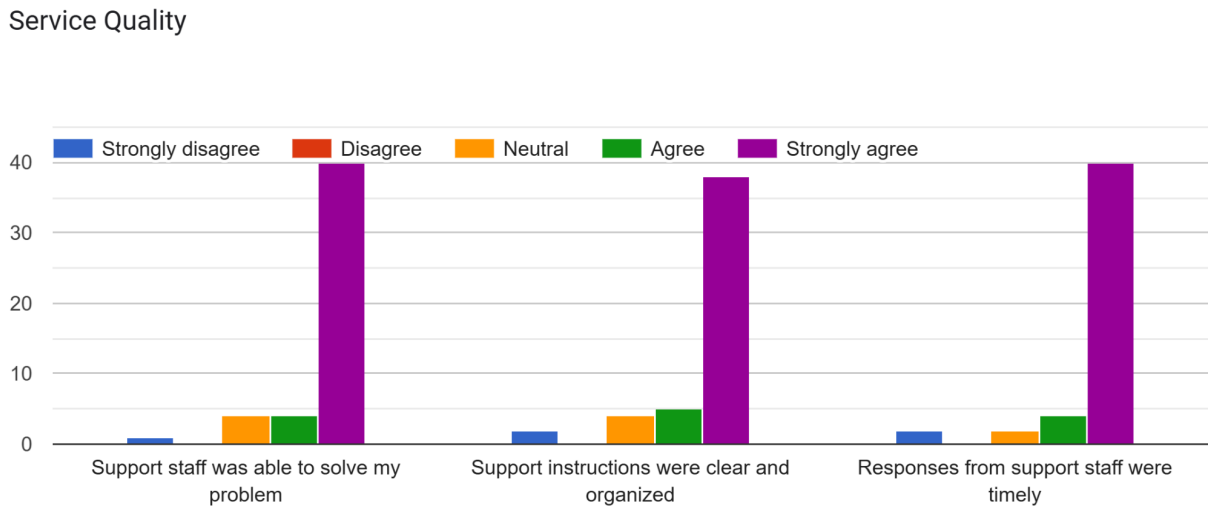

Some of the latest open-ended feedback from the support evaluations includes the following:

- *Tiffany is the most knowledgeable, professional, approachable, fast to answer, very helpful and always spot on. She is just the best. Thank you so much!*
- *Tiffany has done an outstanding job responding to correcting any issues that occur. I truly appreciate all the help she has provided in a very timely manner.*
- *Tiffany is extremely valuable for supporting our efforts to use AWIPS2 in the classroom. She is the most knowledgeable, personable, always helpful and fast to answer. It is a delight to work with her every time. Thank you.*

Software Releases

AWIPS	Date	EDEX	Win	Mac	Linux	VM	python-awips
23.1.1	Sep 10, 2025						x

No major releases have been made since the last status report. All source code is now available. In addition, all python-awips notebooks were updated to work with the latest packages and dealt with all depreciation warnings and errors.

The naming convention the AWIPS Team follows for AWIPS releases tries to follow/reference the NWS. For example, version 23.4.1-1 is based off the NWS 23.4.1 code base where:

23 - is the year the NWS planned on releasing (although this is typically delayed)
4 - is the fiscal quarter the NWS planned on releasing
1 - is the major version
1 - is the subversion that Unidata uses for releases. If there is a preceding “0” that means the release is in beta. Each release made we will increase this number.

Activities Ongoing/In-Progress

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

- creating a docker/podman image of EDEX to review case data (ex. Review radar data)
- looking into users of AWIPS (number of users, University vs public vs government vs private, etc.)
- actively developing new content to incorporate into v23
- maintaining EDEX builds on the Jetstream2 platform, with the help of the Science Gateways team
- responding to all AWIPS support questions from the community and striving to provide realistic solutions in a timely manner
- actively updating and refining our online documentation to be as accurate and useful as possible

Future Activities

Future plans are constantly evolving to meet the needs of our users. With the halt in new capabilities being added by the NWS, Unidata will be looking into adding new products and displays to AWIPS. I am actively participating in conferences, workshops, and virtual message boards (blogs) to expand our user base.

We are looking to put some additional resources towards python-awips to create/update notebooks.

Metrics

~2800 downloads (Sept 2025-Feb 2026)

Prepared *March 2026*

Status Report: IDV

September 2025 - March 2026

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

IDV 7.0u1 released on **September 4, 2025**. With this release, the IDV has successfully migrated to Java 21. This upgrade improves performance, strengthens security, and ensures compatibility with modern Java features, keeping the IDV robust and aligned with current development standards.

IDV System Changes

__IDV Certificates__

Java Windows app and MacOS certificates have been renewed and will be valid until at least May 30, MacOS certificate is valid until 2026). 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 21 AdoptOpenJDK

The Adoptium provides high-performance, cross-platform, open-source prebuilt Java runtime Temurin JDK. The version of the JDK distributed with the IDV is jdk-21.0.6+7. This is a significant step forward in modernizing our software ecosystem, as Java 21 offers many advantages in terms of performance, security, and long-term support. See the Temurin JDK Library for more detailed information.

- IDV employes latest Java3D (1.6.2)

Updated the following Java 3D libraries to the latest versions: j3dcore, j3dutils, vecmath, and jogamp. Replaced the legacy Java Extension mechanism with the explicit classpath-based linking to ensure better modularity, compatibility with modern Java versions, and ease of deployment across platforms.

- Latest netCDF-Java Version (5.9.1)

The version of the netCDF-Java library currently distributed with the IDV is 5.9.0. See the netCDF-Java Library for more detailed information.

IDV Display Changes

__Gemini 2.5 AI process and IDV integrated displays__

We have integrated the Gemini 2.5 AI model into the IDV workflow to enhance data interpretation and visualization through advanced prompt engineering techniques. By carefully designing structured prompts that incorporate meteorological context, variable definitions, and visualization intent, the system significantly improves the model's ability to accurately read and interpret complex IDV-generated images. In addition, the interactive coupling between Gemini 2.5 and IDV displays supports iterative diagnostic analysis, allowing users to refine prompts and obtain progressively more precise insights. This approach not only strengthens automated image understanding but also serves as an effective educational tool, helping students and researchers develop deeper intuition in meteorological diagnostics.

__New IDV default font __

IDV now adopts **Source Code Pro** font family as the new default font in the system, replacing the previous system default font. This update provides significant improvement in the output image and movie. Source Code Pro family provides enhanced readability with balanced character spacing and clean glyph design, offering a more modern and professional

appearance compared to the former default font. This change improves both the aesthetic quality of IDV visualizations and the overall user experience, especially in educational and research environments where clarity and precision are essential.

__3D AIGFS globe view display__

With the AIGFS available in the Unidata TDS server, we created 3D globe view AIGFS display and published bundle in the IDV:

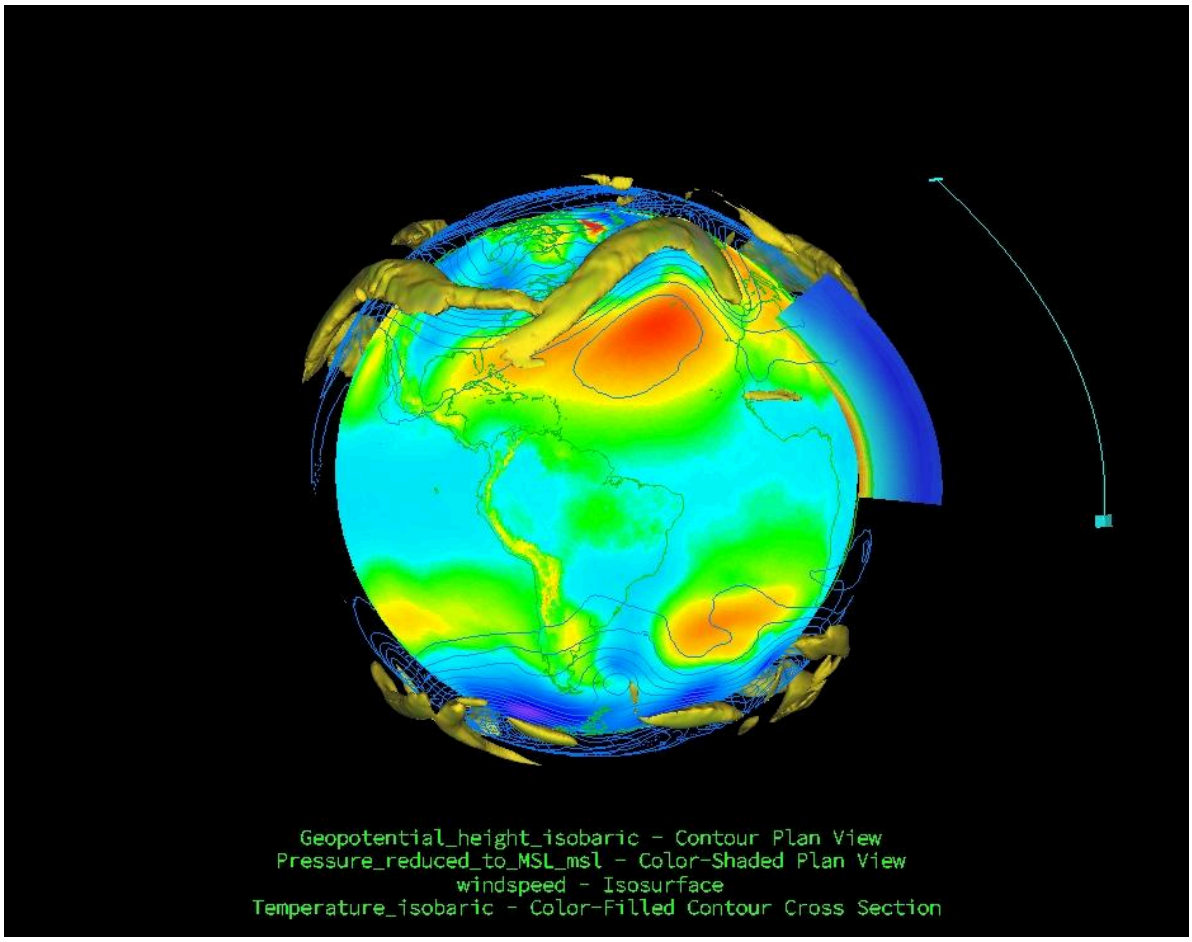

__New Jython formulas: StreamFunction and VelocityPotential__

A new Jython formula has been developed within IDV to compute both the streamfunction and velocity potential from horizontal wind components. This formulation derives the streamfunction from the rotational (non-divergent) component of the flow and the velocity potential from the divergent component, enabling a clear separation of atmospheric motion into its dynamic parts. By solving the corresponding Poisson equations based on the input u and v wind fields, the tool provides users with a more advanced diagnostic capability for analyzing large-scale circulation patterns, divergence, and vorticity. This addition enhances

IDV's support for meteorological education and research by allowing students and researchers to visualize and better understand the fundamental structure of atmospheric flows.

IDV Community Support

Univ of Millersville IDV remote class

We are offering two remote IDV training sessions on October 30 for students at the University of Millersville. These sessions will focus on the basics of IDV as well as accessing and working with satellite data.

Naval Postgraduate School IDV weekly remote meeting

We are hosting weekly remote sessions to support the Lab Manager of the Interactive Digital Environmental Analysis (IDEA) Lab, helping address questions and issues raised by students and professors. This approach represents a new model of IDV support, developed in response to the discontinuation of the traditional three-day, in-person IDV training workshops. To date, we have successfully covered the following topics:

1. Legend and window organization questions for data and how we work with our current data feed structure with overlays.
2. With respect to vertical cross sections, is there a way to have the horizontal analyses read by the transect tool without using an isosurface?
3. If the transect vertical cross-section plot and separate map are in a separate pop-up window from your other horizontal analyses, is there a way to move the vertical cross-section legend items into that pop-up window?
4. For the legend overall, is there a way to reorganize the parameters listed in the legends?
5. Time matching for our data overlays is something we want to understand better, as well.
6. Using the Jython Libraries.
7. Legend modifications
8. How to save bundles properly to avoid loss of work and view/display windows.
9. Building a vector field with formulas VECN of the same specific value.
10. Issues with plotting thickness advection.
11. Building the parallel and perpendicular components for circulation vectors within a vertical cross section.
12. Q-vector plotting.
13. General tips for editing formulas

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).

IDV and RAMADDA Training, Conference Attendance and Presence

Ongoing Activities

We plan to continue the following activities:

__Experimenting IDV 3D output with Gen-AI Gemini__

- Staying updated on new features and capabilities.
- Exploring different ways to integrate Gemini into IDV workflows
- Verifying the information provided by Gemini.
- Learning effective prompting techniques to get the best results.

The future with Gemini holds potential and transformation across geo sciences. We should stay informed with evolving capabilities and embrace the new levels of productivity and creativity.

New Activities

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

We have been in the process of upgrading the version of OPenJDK Java 21. 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.

Prepared *March 2026*

Status Report: Python

September 2025 - March 2026

Ryan May, Drew Camron, Julien Chastang, Ana Espinoza, Tiffany Meyer

Executive Summary

The deadline to apply for travel funding to the 2026 Pythia Hackathon is March 22, 2026. 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. As a result of changing resource availability, we are de-prioritizing synchronous training events to instead prioritize authoring high quality asynchronous online examples and engineering new technical solutions to learning problems in the community, as well as funded efforts around Project Pythia, including the hackathon in June 2206. We presented a workshop in collaboration with Python-AWIPS at the 2026 AMS Student Conference, and are looking at options for AMS 2027. Siphon development is low in immediate priority, but is supported by stable infrastructure and a simmering community engagement. MetPy development continues and 1.8.0 is slated for release late Spring 2026, featuring our C++-based performance enhancements, based on new build, test, and packaging infrastructure. We are exploring new collaborations with other Unidata staff projects, and seeking new project planning and funding opportunities to support MetPy and Siphon development. MetPy's impact on science continues to grow, with 472 theses and peer-reviewed publications mentioning or citing MetPy, including 11 so far in 2026.

Questions for Immediate Committee Feedback

Nothing at this time.

Activities Since the Last Status Report

Python Training

We continue to prioritize a small number of high-impact collaborative synchronous training or development-driven asynchronous training as our resources allow. We welcome committee feedback on this prioritization and any recommendation for focusing our limited training development resources. We recommend committee members consider themselves and their students for participation in the 2026 Pythia Hackathon, June 15–18 (ProjectPythia.org). We thank the community for helping prepare our NSF CSSI proposal on customizable Jupyter data dashboards.

In partnership with the Unidata Science Gateway team, UNC Charlotte, and community members, we submitted a proposal to the NSF to lead development of a Jupyter-based data dashboard prototype and test its impact on student learning and confidence. See the Science Gateway and Cloud Computing status report for more information.

We report on our successful 2026 AMS Student Conference Workshop on AWIPS and MetPy, in partnership with the Unidata AWIPS and Science Gateway teams and past Unidata interns. We delivered to over 60 participants and connected with interested users and prospective student interns. See the AWIPS status report for more information.

We continue to contribute educational content and software infrastructure to Project Pythia, resulting in collaborative updates to Pythia Cookbook usability and growing content on the MetPy Cookbook. We are primary organizers of the upcoming 2026 Hackathon for developing community skills in open science and creating Pythia Cookbooks for open education.

We participate in UCAR/UCP training collaboration with potential to offer and receive support in training development from other organizational experts.

Progress has been made on the following:

- We submitted a proposal to lead collaborative technical development of new opportunity for Python education in the Jupyter ecosystem
- Delivered a cross-program workshop session at the 2026 AMS Student Conference; late Spring we will look to propose one or more workshops for delivery at AMS 2027
- We worked with UCAR/UCP to catalogue our training offerings alongside other programs, and are working towards de-duplication and collaboration
- Unidata continues technical, educational, and community leadership on Project Pythia, including primary involvement in planning and delivery of the annual Pythia Hackathon.

MetPy

Development continues to be driven by requirements for our dedicated awards, in addition to responding to community issues and contributions. MetPy 1.7.1 remains the current release. On-going developments include:

- Using nanobind (instead of pybind11) as a way to write C++ accessible from Python
- Porting calculations from Python to C++
- Reworking MetPy's build/test infrastructure to ensure reliability of shipped packages
- Looking at using Python's **stable ABI** to reduce packaging (and testing) burden

Plans for this work were shared at the 2026 AMS Annual Meeting in Houston, TX. Beyond this, we have continued on-going maintenance for MetPy ensuring that it continues working with the latest versions of packages within the scientific Python ecosystem. We have also invested some in infrastructure around collecting the publication metrics for MetPy, with the goal that this might become useful tooling for other projects.

Progress has been made on the following:

- MetPy 1.8.0 targeting late Spring 2026 release
- Presented on plans for building and shipping C++ code in MetPy 1.8.0 at the 2026 AMS Annual Meeting

- 472 total theses or peer-reviewed publications citing or mentioning MetPy, with 11 so far in 2026.

Siphon

Previously reported Siphon maintenance keeps the CI infrastructure relatively stable, though a few small CI migrations from MetPy are in-waiting. Active development on Siphon is near-zero and there are no new releases to report. Community engagement through Issues and Pull Requests is quiet, but still active.

Proposed Siphon development is currently included in ongoing project & proposal planning in collaboration with THREDDS.

Ongoing Activities

We plan to continue the following activities:

- Engage in support of Project Pythia and adjacent UCAR Python education efforts
- Engage in synchronous Python training when impact- and resource-appropriate
- Deliver enhancements towards the goal of MetPy's dedicated award to address big data challenges
- Maintain Siphon as a tool for remote data access across a variety of services
- Grow and develop MetPy as a community resource for Python in meteorology

New Activities

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

- Deliver 2026 Pythia Hackathon
- Evaluate and submit proposal for synchronous training at 2027 AMS Annual Meeting
- Release MetPy 1.8.0

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

- If grant is awarded, begin technical development on NSF CSSI Data Dashboards
- We are evaluating supporting Python-AWIPS with Python development time
- Deliver one or more high-impact synchronous training events
- Explore and report additional project planning and funding opportunities for Python technical development, internal collaborations, and MetPy Mondays production

Relevant Metrics

- MetPy
    - According to GitHub, 812 repositories and 100 packages depend on MetPy
    - 49 citations/mentions in 2025, 419 total
- Siphon

- According to GitHub, 282 repositories and 36 packages depend on Siphon

---

Prepared *March 2026*

Status Report: IDD, Data Cataloging, and LDM

September 2025 - March 2026

Stonie Cooper, Jennifer Oxelson Ganter, Sean Arms, and Mike Schmidt

Executive Summary

NSF Unidata's LDM developer and IDD maintainer continues to update LDM source code and operating paradigms with ever-changing computing implementations and user requests. The IDD continues to be enhanced with data redundancy and inclusion of new data.

Questions for Immediate Committee Feedback

Not at this time.

Activities Since the Last Status Report

Although not the only technology for distributing data, the IDD is the preferred source for automatically streaming environmental data in real-time. The service-level software for streaming data across the IDD is LDM, and as such, emphasis is placed on maintaining the source code for stability and efficiency.

Progress has been made on the following:

- IDD: Continued monitoring and sourcing redundant data sets.
- Data Cataloging: Implemented RDBMs partitioning to greatly reduce processing time for queries.
- Data Cataloging: Started implementation of the statistics tables.
- Data Cataloging: Opened access for internal-Unidata SQL access for functionality feedback.
- LDM: Final run-time review of the edits to remove run-as-root requirements.
- LDM: Complete rewrite of LDM documentation, modernizing and organizing into chapters. Review copy for internal-Unidata examination in a Wiki instance, includes a 15 section full version, a quick-start guide, and a two-page start guide for system administrators.
- LDM: Started compilation of training materials for LDM.

Dependencies, challenges, problems, and risks include:

- Ever changing landscape of new data types, adjustments from data source paradigms.
- Uncertain scientific development environment on a whole causing work hesitation.

- Reduction in developer and maintainer resources and multitasking across multiple responsibilities.

Ongoing Activities

I plan to continue the following activities:

- Supporting and maintaining LDM and ancillary software.
- Supporting IDD data sourcing with new opportunities and operational observation platforms.
- Implementing data bandwidth metrics to provide real-time data usage accounting.
- Supporting users of the LDM and IDD with quick and professional responses.
- Continuously updating and refining LDM documentation.
- Migration of LDM training to on-demand video/wiki training sessions.

Relevant Metrics

- The LDM software provides the streaming technology for the NSF Unidata IDD network. Metrics on that program are available upon request.
- Number of LDM package downloads immediately following latest version releases.
- Number of support tickets and training requests, views, or downloads.

Prepared *March 2026*

Status Report: THREDDS

September 2025 - March 2026

Sean Arms, Jennifer Oxelson, Ryan May, Ethan Davis

Executive Summary

This THREDDS status report, covering September 2025 to March 2026 focuses on modernization, releases, and community engagement. Key activities include investigating a cross-language architecture for netCDF-Java and the THREDDS Data Server (TDS) to ensure sustainability, alongside a renewed focus on updated and separated documentation for the TDS. The netCDF-Java library was updated with blosc compression for better Zarr support, preparing for the v5.10.0 release. The TDS released version 5.8 with security updates, and v5.9 is imminent with major build and DAP4 improvements. A special TDS war "flavor" now integrates gCDM (gRPC services for the Common Data Model) to facilitate serving data via both Java and Python components. New AI-driven global model outputs (AIGFS, AIGEFS, HGEFS) are now available via Unidata and other community TDSs (notably, the University of Wisconsin's TDS). Future work focuses on netCDF-Java support for cloud-based Feature Collections and Zarr v3, as well as a community survey for the approximately 20 academic institutions running the TDS.

Questions for Immediate Committee Feedback

Question from the previous two status reports: We are actively creating a requirements document for the next generation THREDDS Data Server, and your input is critical. When thinking about accessing data from thredds.ucar.edu, what are your most important features? Does that look different when considering other TDSs? Please reach out to Sean with your thoughts (sarms@ucar.edu)–your input is very important to us!

Activities Since the Last Status Report

General

- We continue investigating options for moving netCDF-Java and the TDS to a more sustainable, cross-language architecture while maintaining as much backwards compatibility as possible.
- Renewing documentation effort on netCDF-Java and TDS, which is in alignment with some of the the findings from the NSF Unidata Community Assessment survey.

NetCDF-Java

- Migration away from Sonatype NEXUS artifacts server–appears to have been a seamless transition for our Java community (for those only familiar with Python, think of it as a Unidata specific PyPI, but for Java artifacts).

- Working on the v5.10.0 release of netCDF-Java
- Support for blosc compression has been added. This significantly improves support for Zarr in netCDF-Java.

TDS

- Version 5.8 was released in October of 2025. This contained a number of bug fixes, security updates, and improvements. See also: https://github.com/Unidata/tds/releases/tag/v5.8.
- We are working on v5.9 of the TDS and anticipate a release later this spring. The build system received a major upgrade, which has significantly improved the thredds.war application and greatly reduced the number of warning messages on startup, and many improvements to DAP4 have been made. Additionally, updates to the netCDF-Java library will improve the functionality and stability of the TDS.
- gCDM, which was the cornerstone of our icechunk serving capability demonstration at the Fall meeting, has been integrated into a special “flavor” of the thredds.war package. Progress is being made on the python based gCDM-Icechunk package, which completes the pipeline for serving IceChunk data through the TDS.
- The Unidata organization on dockerhub is now part of the Docker-Sponsored Open Source (DSOS) program. Among other benefits, this means DockerHub’s rate limits do not apply when pulling public images from the Unidata organization.

Ongoing Activities

Server management

- Unidata hosts and maintains the following THREDDS servers:
    - thredds.ucar.edu always runs the latest stable release of the TDS (unless a quick security update is required).
    - thredds-test.unidata.ucar.edu always runs the latest development version.
    - thredds-dev.unidata.ucar.edu will be the testbed for the new cross-language TDS. It is not expected that this server will host the same set of data as thredds.ucar.edu or thredds-test.unidata.ucar.edu, although there may be some overlap.
        - This is where you can see the Java - Python gCDM connection for serving IceChunk data in action, as well as serving data from other cloud native storage locations.
    - Additionally, the Cloud Activities group manages cloud-hosted TDS instances (see cloud activities report for details).

New Datasets

Output from the AI-driven global weather prediction suite:

- AIGFS (Artificial Intelligence Global Forecast System)
- AIGEFS (Artificial Intelligence Global Ensemble Forecast System)

- HGEFS (Hybrid-GEFS - a hybrid ensemble of AIGEFS and Global Ensemble Forecast System output)

are now available on the Unidata TDS. See the announcement on the Unidata blog for more details.

Development

- netCDF-Java
  - We're actively working to extend the ability for netCDF-Java to support cloud-based Feature Collections in the TDS, including GRIB featureCollections.
  - While we continue to improve support for Zarr v2, we're beginning to look into consolidated metadata, as well as Zarr v3 support.
- TDS
  - We are preparing a survey for the roughly 20 academic institutions running the TDS. Our goal is to gain insight in their specific needs and use cases when it comes to running the TDS at their institution.
  - As reported during the last meeting, in alignment with some of the findings from the NSF Unidata Community Assessment survey, we are renewing documentation efforts for the TDS. Currently, the "user guide" is more aligned with a "server administrators" guide and tutorial. We are actively working to split and enhance the user guide into three separate sets of documentation:
    - A new user guide, focused on non-server specific end-user access to the TDS (how to use various services, API documentation, lists of tools known to interact with the TDS)
    - A quickstart guide, designed to get administrators up-and-running with a TDS quickly
    - An admin guide, complete with a full tutorial and full documentation on how to run a TDS

The following active proposals directly involve THREDDS work:

- None at this time.

New Activities

As development progresses on the technology to bridge the Java and python ecosystems (in support of a more sustainable TDS tech stack), we will continue to seek out datasets to exercise this new functionality. Along the lines of the icechunk store of ERA5 analysis, what other datasets would you like to see us target for development? We are also preparing a survey for the ~20 academic institutions running a TDS; this survey will largely inform our future development paths for the TDS.

Relevant Metrics

THREDDS Startup Metrics

	2025-09 — 2026-02	2014-08 — 2026-09
TDS Startup (unique IP address count)	3,566	50,417

	Total Servers	Information page updated
Publicly Accessible[1] TDS count	97	51

Over the past 5 months, **3,566** unique IPs started up the TDS (September 2025 through February 2026). Since we've started tracking these metrics (v4.5.3, August 26th, 2014), we've seen the TDS startup from **50,414** unique IP addresses. There are currently **97** publically accessible TDSs running "in the wild". Of the **97** publically accessible servers, **51** 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 v5.4 is the dominant specific version running in the wild (although more have migrated to TDS v5.6), and we are seeing a shift to the latest (v5.8, as of this report), which was released in September 17h, 2025.

[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.

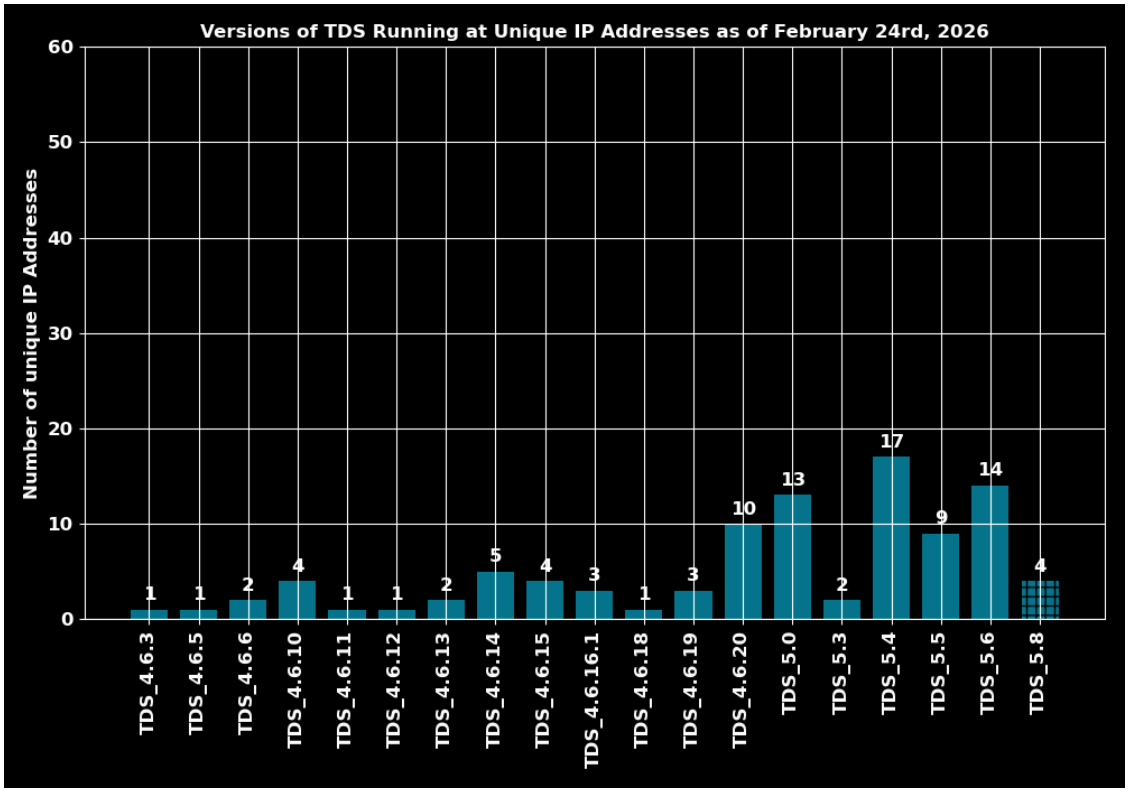

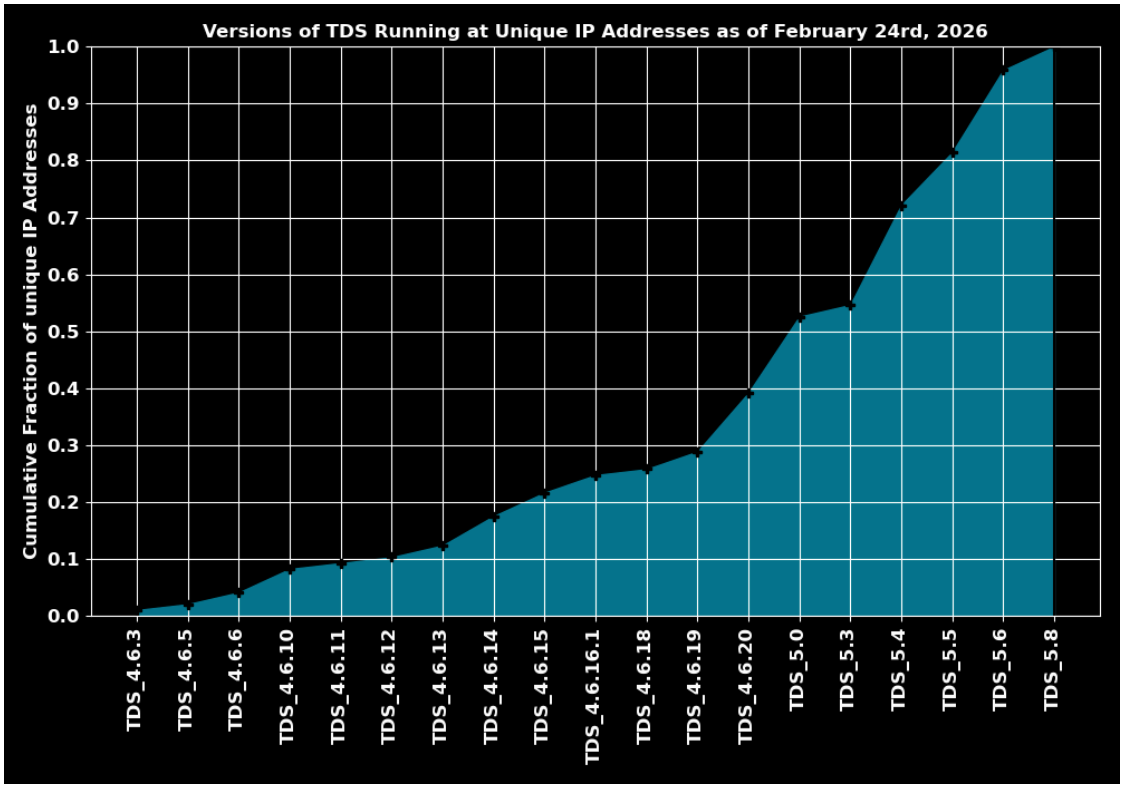

While mining startup metrics has proven useful, the startup logs only capture part of the picture, as checking for updates on startup is optional. In early June of 2025, UPC

staff took a deep dive into TDSs running in the wild using a variety of techniques and found a total of 108 publicly accessible TDSs that appear to be intended for external use (as opposed to the 53 identified by mining startup logs alone). *While this was outlined in the last update, it is repeated here to make sure this information gets out to the new members of the committees.* The breakdown of sites by high level entity type is as follows:

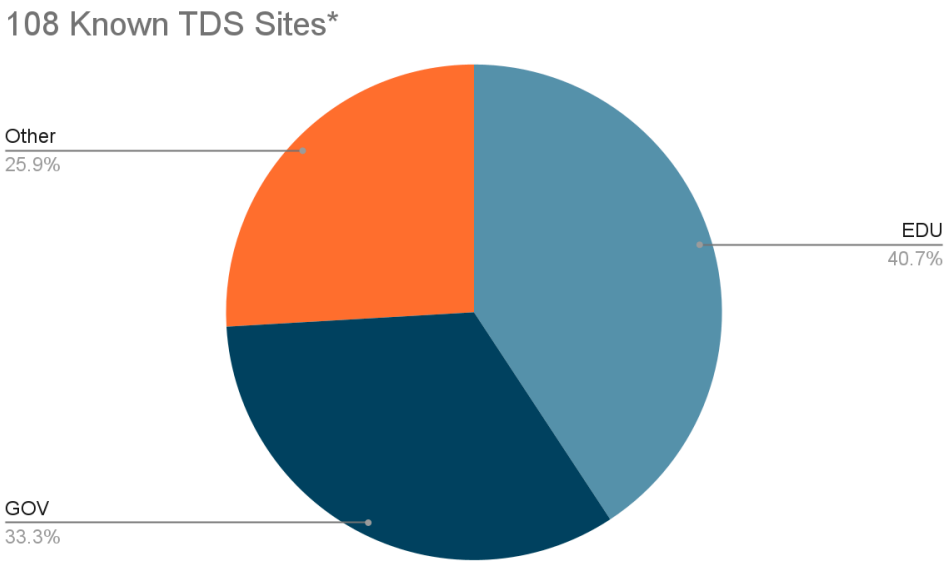

These sites are roughly split between US and NON-US institutions:

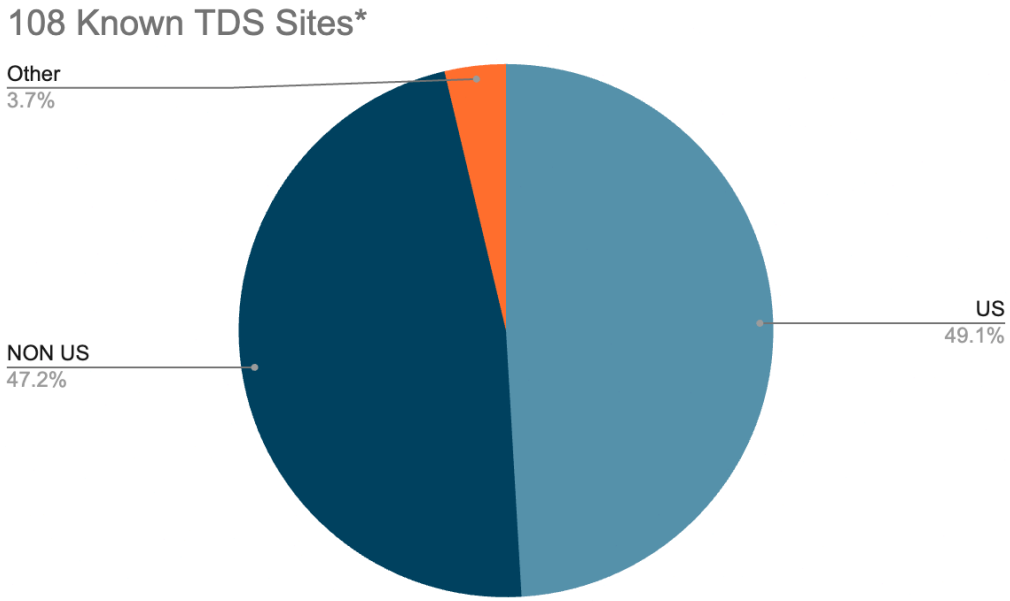

20 US Academic institutions currently operate one or more publicly accessible TDSs:

Prepared *March 2026*

Status Report: Science Gateway and Cloud Computing Activities

September 2025 - March 2026

Sean Arms, Julien Chastang, Ethan Davis, Doug Dirks,
Ana Espinoza, Ward Fisher, Ryan May, Tiffany Meyer, Jennifer Oxelson Ganter, Mike Schmidt,
Tanya Vance

Executive Summary

- Secured 5.1M CPU SUs, 467k GPU SUs and 40TB of storage on Jetstream2 effective October 1, 2025.
- Fully transitioned PyAOS JupyteHub deployments to OpenStack Magnum autoscaling replacing older Kubespray/Terraform/Ansible workflow.
- Implemented shared user storage with per-user quotas eliminating volume attachment errors
- Deployed browser-based AWIPS CAVE and IDV desktops within JupyterLab enabling notebook to GUI workflows.
- Served 334 students at 13 institutions (Fall 2025 - Spring 2026)
- Deployed FluentBit centralized logging for all JupyterHub clusters to meet security auditing requirements.
- Launched CILogon authentication on jupyterhub.unidata.ucar.edu allowing access from institutions alongside additional choices.
- Developed prototype “on-demand notebook gallery” integrating nbgallery with JupyterHub for frictionless notebook launch.
- Began deploying conda-store to enable user-defined, shareable, reproducible environments on JupyterHub deployments
- Modernized Gateway web presence via migration to UCAR Drupal platform.
- Began splitting monolithic science-gateway GitHub repository into modular, DevOps, VM stack, and Jupyter configuration repositories.
- Released thredds-docker 5.8 aligned with TDS 5.8 release.
- Continued AWS-hosted NEXRAD Level II/III data via THREDDS.
- Hosted AWIPS EDEX production and development VMs on Jetstream2 with up to 300 daily users.
- Submitted CSS “Elements” proposal ($600k/3 yrs) for composable web-based dashboard.
- Participated in CSSI “Framework Implementations” proposal to port LROSE Gateway to NSF NCAR GDEX
- Began investigation of IceChunk MRMS object-store on Jetstream2

Questions for Immediate Committee Feedback

None at this time

Activities Since the Last Status Report

Jetstream2 2025-26 NSF ACCESS Grant Approval

NSF Unidata successfully secured a grant for 5,104,375.0 CPU SUs, 467,000.0 GPU SUs and 40TB of Indiana Jetstream2 storage from the NSF ACCESS program for the current Jetstream2 allocation cycle. This allocation, effective October 1, 2025, ensures continued access to essential servers such as AWIPS EDEX, JupyterHub, THREDDS, RAMADDA, LDM/IDD nodes, and Science Gateway resources. The requested and ultimately awarded resources are lower than those received for the 2024-25 cycle due to efficiency gains from Magnum auto-scaling cluster deployments and reduced GPU usage following the departure of Thomas Martin, which resulted in a lack of in-house AI/ML expertise. In the future, we hope to request an amount below the 5,000,000 CPU SUs limit which will trigger a simpler allocation request process.

CSSI Grant Proposal Submissions

Alongside several software engineers across the Unidata Program Center and Dr. Casey Davenport of UNC-Charlotte, the Science Gateway team submitted a proposal to the National Science Foundation's Cyberinfrastructure for Sustained Scientific Innovation (CSSI) NSF 22-632 solicitation (Elements track). If awarded, NSF Unidata will be funded for up to $600,000 over 3 years to develop a user-composable, web-based data dashboard framework that can act as a springboard from rapid data exploration to in-depth, code-driven data explanation.

In addition, NSF Unidata Science Gateway staff were considered in a proposal for the same solicitation (Framework Implementations track) submitted by a collaboration between UCAR/NSF NCAR and CSU scientists and engineers to, among other activities, port the LROSE Science Gateway to NSF NCAR compute resources, connecting gateway visitors to the vast and desirable data sets found on NSF NCAR's Geoscience Data EXchange (GDEX) research data platform. If awarded, NSF Unidata Science Gateway staff will be funded for 3-5 years to participate in LROSE Science Gateway and GDEX activities.

Science Gateway Re-imagined (SGRI) Activities

"On-Demand Notebook Gallery": nbgallery and Jupyterhub Integration Prototype

We have developed a prototype "on-demand notebook" capability using the open-source nbgallery project. The prototype currently hosted at https://nbgallery.ees220002.projects.jetstream-cloud.org stores curated Jupyter notebooks that authenticated users can browse and launch at https://jupyterhub.unidata.ucar.edu. The goal is to incorporate this functionality into the Science Gateway as part of the SGRI effort allowing users to launch notebooks without preliminary setup and minimal friction.

We used agentic AI tools to assist in developing the customization and integration of these technologies to securely launch notebooks from nbgallery to JupyterHub user sessions. The prototype is functional, and we are evaluating the agentic approach for feasibility, security, and long-term maintainability before considering a broader deployment.

JupyterHub User Managed Environments & Environment Sharing

Today, if a JupyterHub user on the Science Gateway needs an update to their conda environment they must contact the Gateway team to ensure the update gets rolled out to all JupyterHub users. This is necessitated by the containerization approach to our JupyterHub deployments, but can be slow and will become infeasible as our user base grows. We have deployed an instance of conda-store, a conda-environment web service, on Jetstream2 for development of the techniques necessary to connect conda-store with the rest of the Gateway.

In its final form, users will be able to login to conda-store using the same credentials used for JupyterHub. From the conda-store UI, they will then define conda environments that are built on the server and distributed to all users via a shared file system approach. In addition, visitors to the conda-store website can download build artifacts, such as conda lockfiles, environment files with pinned package versions, and `.tar.gz` archives containing the environment. This enables consistent reproducibility of compute environments for sharing with colleagues or using the environment on other cyberinfrastructure.

Development is on-going as we explore how to:

- Connect conda-store with CILogon (see the CILogon Authentication section below) and jupyterhub.unidata.ucar.edu
- Implement a JS2 native “Manila share” “File System as a Service” to host environments across all JupyterHubs of the Gateway
- Robustly configure conda-store to allow proper authorization and permissions for users to create, read, edit, and delete environments

Science Gateway Website Modernization

Migrated away from the old https://science-gateway.unidata.ucar.edu/ science gateway landing page to https://www.unidata.ucar.edu/science-gateway which makes use of UCAR’s Drupal supported platform and as part of their web modernization initiative.

Education Hub: NSF Unidata and COMET Collaboration

Some of the learning resources found in the Science Gateway’s Education Hub are accessed via NSF Unidata’s https://eLearning.unidata.ucar.edu website. This website runs on the Moodle Learning Management System (LMS) hosted on NSF Unidata resources. Work is underway, spearheaded by Tanya Vance of the Community Services group, to integrate our eLearning resources with the popular MetEd system, running on its own LMS managed by the COMET program.

Science Gateway Repository Migration Modernization

For several years, https://github.com/Unidata/science-gateway has served as the open-source repository for Science Gateway code and configurations. Over time, this repository grew into a large, monolithic codebase encompassing scripts related to Jetstream2 OpenStack VMs, JupyterHub Kubernetes configurations for semester-long classes and workshops, and deployment artifacts related to the various services we have running on Jestream2 (e.g., THREDDS, AWIPS EDEX). Over time, this repository has become unwieldy encompassing too many disparate concerns.

We are now migrating to a more modular repository plan with three repositories:

- https://github.com/Unidata/usg-vm-stacks - Docker configurations related to Science Gateway services (e.g., THREDDS, AWIPS EDEX).
- https://github.com/Unidata/usg-devops - deployment automation scripts related to Jetstream2 VMs and clusters
- https://github.com/Unidata/usg-jupyter-configs - JupyterHub configurations for classes and workshops
- https://github.com/Unidata/science-gateway - The current, monolithic repository will be trimmed down and used to host user-facing materials, such as Gateway user documentation

This new separation of concerns will enable us to more easily share our work with the wider Jetstream2, science gateways, and earth systems science communities, as well as automate more of our development to operations (DevOps) workflows for increased developer efficiency.

CILogon Authentication Implementation and Institutional Access

We implemented CILogon authentication on https://jupyterhub.unidata.ucar.edu enabling users to login using credentials from their home institutions. While the current GitHub OAuth mechanism always works, we hope this new access workflow will ultimately be part of an academically focused, federated, single sign-on strategy for Science Gateway resources.

In addition to institution based logins, authentication via NSF ACCESS, ORCID and GitHub are also available. These providers are currently managed through allow-listing to ensure appropriate use of these resources.

JupyterHub Activities

Transition to Shared User Storage With Per-User Quotas

Previously, each JupyterHub user was assigned an individual persistent storage volume. While this model functioned adequately in most cases, it occasionally led to failures in volume attachments resulting in “stuck” accounts often during active teaching periods requiring intervention from Unidata and Jetstream2 staff.

We have transitioned to a shared storage architecture modeled after the 2i2c approach. In this configuration, all users access a single, continuously mounted storage with per-user disk quotas. Because the storage is always attached to the JupyterHub environment, this design eliminates many of the disk attachment failures that previously caused teaching disruptions.

This architecture is more robust and also positions us to implement more systematic backup strategies in the future as user data now resides in a central place instead of across numerous independent volumes.

PyAOS JupyterHub and Magnum Autoscaling on Jetstream2

For Spring 2026, we completely transitioned to an OpenStack Magnum-based deployment workflow for our customized PyAOS JupyterHub environments on Jetstream2. Previously, clusters were provisioned to accommodate peak anticipated demand often leading to underutilized resources during quieter periods (e.g, outside classroom time). With Magnum's autoscaling capabilities, clusters can now expand and contract in response to actual usage. Our current usage metrics indicate that our service-unit consumption is tracking with projected estimates. Using Nushell, we also refactored our deployment workflow to make Magnum-based clusters modular, reproducible, standardized and faster for provisioning. In essence, Magnum and Nushell has completely replaced the previous Kubespray/Terraform/Ansible workflow.

JupyterHub Virtual Desktop Technology for IDV and CAVE

For Spring 2026, AWIPS CAVE and IDV are now available across all our JupyterHub servers through browser-based streaming Desktop technology. Users access these applications via the "Unidata Desktop" button in the JupyterLab interface. These new features provide access to mature GUI-based applications along Jupyter notebooks in one integrated environment requiring no local software installation. As an example, during fall semester of 2025, Professor Brian Mapes (University of Miami) created a workflow retrieving GFS data from a Unidata THREDDS server, used the Python windspharm library to compute Helmholtz Decomposition values and wrote the results to a netCDF file. This file was then read into the Unidata IDV, and satellite IR imagery data was overlaid on the data all within the same Jupyter session. The result is a unified environment where students can perform computational analysis and immediately explore the results using the IDV without requiring any local software installation.

Kubernetes Fluent Bit Logging for PyAOS JupyterHub Clusters

To comply with security best practices, we store, compress, and archive all cluster logs generated through a JupyterHub's lifetime. This is done using the fluentbit technology, configured such that each cluster forwards generated logs to a central logging server also found on Jetstream2. Logs can be backed up for storage at the Unidata Program Center.

JupyterHub Servers for Workshops, Fall and Spring Semesters

NSF 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

a steady pace. NSF 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 2025 encompassing the length of this status report , 334 students at 13 academic institutions and various workshops have used NSF Unidata JupyterHub servers running on Jetstream2.

Notably, we provided JupyterHub resources to:

- New Institutions:
    - College of Dupage
    - University of Georgia

LROSE Collaboration between Colorado State University (CSU) and NSF NCAR EOL

The NSF Unidata Science Gateway team continues its collaboration with Professor Michael Bell's group at CSU and NSF NCAR's Earth Observing Laboratory (EOL) to enhance their science gateway for radar meteorology. This gateway features a JupyterHub environment integrated with LROSE (Lidar Radar Open Software Environment) to support advanced data analysis and visualization.

Current efforts are centered around assisting Brenda Javornik (EOL) and Jen DeHart (CSU) with Jestream2 planned usage and an NSF ACCESS request renewal to support the LROSE Science Gateway.

Unidata Participating in the Docker-Sponsored Open Source (DSOS) Program

The Unidata organization on DockerHub is now part of the Docker-Sponsored Open Source (DSOS) program. Among other benefits, this means DockerHub's rate limits do not apply when pulling public images from the Unidata organization.

Ongoing Activities

NOAA Big Data Program

- NSF Unidata continues to manage the NEXRAD level 2 archive in Amazon S3, ensuring that realtime data are successfully delivered to the noaa-nexrad-level2 bucket. LDM is employed to deliver these data.
- Public Bucket for level II NEXRAD: https://s3.amazonaws.com/noaa-nexrad-level2/index.html
- NSF Unidata also continues to deliver NEXRAD level 3 products to the NSF Unidata-nexrad-level3 bucket, part of the AWS public datasets program.
- TDS on Jetstream2 for level II NEXRAD: https://tds-nexrad.scigw.unidata.ucar.edu/thredds/catalog/catalog.html. Recent improvements to this server have been made with regards to caching strategy, which

has improved the overall reliability of the service.

Jetstream2 Allocation Management and Collaboration

NSF Unidata staff continue to actively manage our Jetstream2 allocation and collaborate with the Jetstream2 team. We closely monitor our allocation and decommission outdated resources to prevent allocation exhaustion.

Jetstream2 and Science Gateway Security

We continually work with NSF 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.

Collaboration with Andrea Zonca and Julian Pistorius on Jetstream2

Now that we have completely transitioned to OpenStack Magnum for JupyterHub cluster deployments, we are working with Jetstream2 to ensure stable operations in this area. While the autoscaling model has improved efficiency, we have encountered several operational issues:

- **Cluster initialization failures**: On occasion, a Magnum cluster becomes stuck during initial deployment requiring the cluster creation process to be restarted in an otherwise nearly completely automated workflow.
- **Load balancer failures**: In rare cases, load balancers have become unresponsive, rendering the cluster inoperable. In one instance, this required launching an entirely new cluster and migrating user data.

Aside from these issues, we worked with Andrea Zonca on the shared user 2i2c storage solution described earlier so that the Jetstream2 team can document and promote this more robust deployment strategy.

Frequent collaboration continues over the Jetstream2 Matrix channel, where we communicate with Andrea and Julian when issues arise. This long-standing partnership has been invaluable, providing ongoing expertise and guidance that has yielded many benefits for both the Science Gateway and the Jetstream2 projects.

Docker Containerization of NSF Unidata Technology

We continue to employ Docker container technology to streamline building, deploying, and running NSF Unidata technology offerings in cloud-based environments. Specifically, we are refining and improving Docker images for the LDM, RAMADDA, THREDDS (TDS), and the THREDDS Data Manager (TDM). Most recently, we released thredds-docker 5.8 in conjunction with the 5.8 release of the TDS. In addition, we also maintain a security-hardened NSF Unidata Tomcat container inherited by the RAMADDA and THREDDS 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

NSF 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. We've had upwards of 300 clients connecting to EDEX in a single day. 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.

Because we are needing to spin up new machines fairly often, we have simplified and streamlined this process by creating custom Rocky 8 images that can be used for deployment on Jetstream2. 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 2 sets of servers as described above: one set has been running our v23 production software, another running v23 development software. We may be looking to add an additional 2 back in the mix (running Rocky 8) for future development and beta builds. Having backup/development servers allows us to be able to patch, maintain, and develop our servers while still having a functional server for our users and to minimize any down time.

Nexrad AWS THREDDS Server on Jetstream2 Cloud

As part of the NOAA Big Data Project, NSF Unidata maintains a THREDDS data server on the Jetstream2 cloud serving Nexrad data from Amazon S3. This TDS server leverages Internet2 high bandwidth capability for serving the radar data from Amazon S3 data holdings. TDS team and science gateway staff collaborate to maintain this server.

Presentations/Publications/Posters

```
@conference{Chastang2026a,
  title        = {{Unidata Science Gateway: Integrating Atmospheric Science Desktop
Applications with Modern Code-Driven Workflows in the Cloud}},
  author       = {Chastang, Julien and Espinoza, Ana},
  year         = 2026,
```

```
  month        = {Jan 25--29},
  booktitle    = {Proceedings, 42nd Conference on Environmental Information Processing
Technologies, 106th AMS Annual Meeting},
  address      = {Houston, Texas, USA},
  doi          = {10.6084/m9.figshare.31325752.v1},
  url          = {https://ams.confex.com/ams/106ANNUAL/meetingapp.cgi/Paper/477111},
  note         = {See also https://science-gateway.unidata.ucar.edu},
  organization = AMS
}
@conference{Chastang2025b,
  title        = {{Integrating Scientific GUI Applications into Autoscaled Science Gateway
Environments}},
  author       = {Chastang, Julien and Espinoza, Ana},
  year         = 2025,
  month        = {Oct 28--30},
  booktitle    = {Proceedings of the Gateways 2025 Conference},
  address      = {Green Bay, Wisconsin, USA},
  doi          = {10.6084/m9.figshare.30566273.v1},
  url          = {https://doi.org/10.6084/m9.figshare.30566273.v1}
}
```

New Activities

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

Forthcoming conference participation

- Over the next three months, we plan to develop a workflow for hosting an IceChunk store of MRMS analysis on Jetstream's Object Store.

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

- We are continuing to test Magnum autoscaling clusters in combination with virtual desktop technology to support traditional applications such as IDV and AWIPS CAVE, with multiple deployments currently underway for spring 2026.
- Prof. Ronald Stenz (College of DuPage) requested computing support for a Spring 2026 mesoscale class, including a PyAOS JupyterHub environment with the capability to run the CM1 model. We plan to explore this request further, as it represents an interesting use case.
- We plan on exploring integrating a small podman EDEX container into a JupyterHub with instructions on how to load archived data and visualize via CAVE.

Relevant Metrics

Spring 2026 / Fall 2025 JupyterHub Servers

Since spring of 2020, NSF Unidata has provided access to JupyterHub scientific computing resources to nearly 2600 researchers, educators, and students (including a few NSF REU students) at 29 universities, workshops (regional, AMS, online), and the UCAR SOARS program. Below are the latest metrics (institution, number of active users, point of contact) since the last status report.

Fall 2025		
UNC Charlotte	1	Wesley Taylor
Florida Institute of Technology	32	Steve Lazarus
Millersville	34	Sepi
University of Miami	11	Brian Mapes
Southern Arkansas University	41	Keith Maull
Seoul National University	15	Duseong Jo
University of Wisconsin	29	Hannah Zanowski
University of Wisconsin Dask	29	Hannah Zanowski
Vermont State University	11	Andrew Westgate
Spring 2026		
AMS 2026 Student workshop	51	According to login spreadsheet
Seoul National University	1	Duseong Jo
Florida Institute of Technology	4	Milla Costa
SUNY Oswego	2	Scott Steiger
Millerville	21	
College of Dupage	0	
University of Georgia	23	
Florida State University	23	Christopher Homles
Vermont State University	6	
Virginia Tech	0	

Note: Some entries in the table above indicate zero or one user. These are recently launched Hubs and the instructors have not yet had a chance to complete their setup and launch the coursework for students to access.

Jetstream2 Allocation Usage Overview

For the period between October 1, 2024 and September 30, 2025, NSF Unidata was granted

an allocation on the NSF Jetstream2 (JS2) cloud valued at 8.6M+ CPU Service Units and 350k GPU Service Units (SUs), the unit of "currency" on JS2. NSF Unidata staff has been proactive in ensuring JS2 resources are being used effectively in a non-wasteful manner by automating SU usage data collection through interactions with the JS2 API. This data is extrapolated forward in time to predict future SU usage, allowing us to make meaningful decisions about the gateway's capabilities.

As of February 27, 2026, our allocation usage is as follows:

Resource	SUs Allocated	SUs Used	SUs Remaining	Daily Usage Rate (7 day avg)	Over/Under-Budget
CPU	5.10M SU	2.39M SU	2.71M SU	17.2 k SU/day	Over
GPU	467k SU	16.0k SU	451k SU	0 SU/day	Under

While our CPU SUs are over-budget at our current usage rate, this is not entirely unexpected or cause for concern. High gateway usage, and thus high SU usage is expected to coincide with the academic semesters. We expect the SU usage rate to fall at the end of the spring 2026 semester when JupyterHub clusters are torn down, bringing us back under budget or breaking even.

As this is our first year with Magnum/auto-scaling only clusters on Jetstream2, our estimates for SU usage may need refinement. In the past, Jetstream2 has been amenable to supplemental SU grants in the event more SUs are needed.

Github Statistics*

Repository	Watches	Stars	Forks	Open Issues	Closed Issues	Open PRs	Closed PRs
science-gateway	5 (-1)	20 (+1)	13	6	167	13 (-4)	899 (+39)
tomcat-docker	9 (-1)	66 (-1)	70 (-1)	0	42	1	97
thredds-docker	12 (-1)	40	33(+3)	3 (-1)	127 (+2)	0	188 (+5)
ramadda-docker	2 (-1)	0	2	1	10	0	38
ldm-docker	8	13 (+1)	15	1 (-1)	42 (+2)	0	78 (+8)
tdm-docker	3 (-1)	4	7	0	10	0	31 (+2)

* Numbers in parentheses denote change from last stat report

Prepared *March 2026*

Status Report: Data Standards and Technical Engagement

September 2025 - March 2026

Ethan Davis, Ward Fisher, Sean Arms, Stonie Cooper, and Ryan May

Executive Summary

Engage with federal science agencies, international standards bodies, and other communities focused on data and technology including NASA, NOAA, USGS, World Meteorological Organization (WMO), Open Geospatial Consortium (OGC), Earth System Information Partners (ESIP), CF Conventions for netCDF community, OPeNDAP, and the Zarr and GeoZarr community.

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.

Questions for Immediate Committee Feedback

No questions at this time.

Activities Since the Last Status Report

WMO WIS 2.0

For the last several years, Unidata has been tracking the development of the WMO Information System 2.0 (WIS2) and discussing, internally and with WIS2 developers and data managers, how it might interact with Unidata technologies like LDM and THREDDS. Moving forward (2-6 years) WIS2 will be the framework for WMO data sharing, eventually replacing the GTS. WIS2 is built on modern messaging standards (e.g., MQTT) and OGC metadata, catalog, and data standards.

The LDM team has experimented with running a WIS 2.0 Node, using the wis2box package (docs)(GH), to better understand WIS 2.0 technologies. Further investigation will be needed.

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.

The Zarr v3 specification continues to evolve. The Zarr developers have been struggling with community implementations outpacing the specification process, resulting in a growing divergence between community-supported features and the features outlined in the specification. There exists a reality in which the netCDF team must consider community supported features which deviate from the spec, in order to maintain broad compatibility, although we are hoping to minimize this.

Progress has been made on the following:

- Members of Unidata's netCDF teams 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.

ESIP Activities

During the January 2026 ESIP meeting, two efforts were spun up that are being tracked and contributed to by Unidata staff:

- Development of the process for advancing ACDD (Attribute Convention for Data Discovery) which has been at v 1.3 since 2015.
- Discussions around developing a standard for chunk manifests (e.g., Kerchunk, Icechunk, DMR++)

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 2025 CF workshop (website) was held virtually 22-25 September 2025.
- Ethan Davis continues serving as chair of the CF Governance Panel.
- Members of the CF Governance Panel and CF Committees have been participating in talks with the GeoZarr working group chairs.

OGC activities

UCAR hosted the Oct 2025 OGC Member Meeting, 28-31 October 2025 at UCAR CG1.

We continue to follow and participate in the OGC MetOcean working group as well as the working groups for standards being used by the WMO WIS 2.0.

OPeNDAP activities

Unidata and the OPeNDAP group have been working together to strengthen the DAP4 specification. Part of this effort has been to bring the behaviours of the Hyrax and TDS implementations of DAP4 into better alignment using PyDAP as an independent client. These efforts have been fruitful in identifying various bugs across all three platforms involved. The OPeNDAP group is leading discussion efforts at ESIP regarding the development of a standardized chunk manifest.

Ongoing Activities

We plan to continue the following activities:

- Track and engage in WMO data standards efforts
- Continue efforts to update and reorganize the NetCDF User's Guide (NUG)
- Continue conversations with CF, WMO, ECMWF, and others on the development of mappings between CF Standard Names and GRIB/BUFR variable names.
- Represent Unidata in Earth System Information Partners
- Represent UCAR and Unidata in OGC and various OGC working groups
- Organize regular meetings of the OGC netCDF SWG.

Prepared *March 2026*

Status Report: netCDF

September 2025 - March 2026

Ward Fisher , *Ethan Davis* , *Sean Arms*

Project Summary: Green

Executive Summary

The netCDF team continues to work towards maintaining the sustainability and viability of the netCDF libraries. While facing challenges when prioritizing work against the resources available, we are fortunate to have an engaged community of users and developers.

The status of the netCDF project in summary:

> **NetCDF is healthy and remains viable, thanks to the engagement and support of our community.**

Our efforts to serve the community are reciprocated, through high levels of engagement and contributions, for which we are immensely grateful. The netCDF team lacks the resources to quickly evaluate every potentially useful emergent technology or address every bug report, and we must therefore triage based on what best serves our communities interests at large. We continue to advocate for our community through participation in external data-oriented/focused groups.

Questions for the Committee

How can we encourage additional community engagement, from students and/or faculty? We benefit greatly from the involvement of our community, making netCDF truly a collaborative effort. How can we encourage/expand this collaboration? What makes it rewarding to engage with the netCDF developers and the netCDF project. How should we prioritize our limited resources?

NetCDF Project Status

Team Status

The active team is composed of the C/Fortran/C++ team lead, Ward Fisher, Sean Arms as the lead developer on netCDF-Java, and our community of developers. Work continues apace, with a strong focus on community building and bug fixing.

Status of Community Relationships

NetCDF User and Developer Community

NetCDF continues to enjoy a high amount of community engagement, for which we are very grateful. The primary avenues of engagement with the netCDF community are as follows:

- Committee Meetings (always a pleasure, you're all great).
- Github (issues, conversations)
- Email - Direct, or through the NSF Unidata eSupport system
- Professional events - AGU, RMCC HPC Symposium, SEA ISS Conference
- Invited Talks

Broader Community Engagement

The netCDF team continues to represent the netCDF community in the following areas:

- Zarr Community Meetings: The Zarr implementation council has reorganized, and while NSF Unidata no longer represents a seat on a committee, we continue to represent our community via engagement in the community meetings and through the established relationship with the developers behind Zarr.
- We have re-established connections with the HDF group and the current hdf5 development team. This has led to a number of very good conversations and will hopefully continue to bear fruit in the form of reduced friction between new versions of `libhdf5` and `libnetcdf`. We've already seen this around breaking issues in `libhdf5 2.0.0`, and a commitment from the HDF Group to avoid making the same mistake.
- Working, ongoing collaboration with developers at the German Climate Computing Center (DKRZ) to enhance `ncZarr`.
- Various one-off symposia, conferences, workshops, etc, focused on data formats and data standards (see Data Standards for more information).

Short-Term Priorities

With the release of netCDF `netCDF-C v4.10.0`, we have turned to revitalizing our infrastructure, so that the work we do can be accomplished more effectively and efficiently. This has included revitalization of our `docker`-based regression testing, `docker.unidata.ucar.edu/nctests` (image) (configuration files here: https://github.com/Unidata/docker-nctests). This will help improve the speed of our development cycle, and help us avoid situations where we have to fix interface-library-breaking bugs that inadvertently made their way into a release. Short-term priorities for netCDF-Java can be found in the THREDDS Status Report.

Challenges

Lack of resources refers to "Not having enough developer hours to address all of the issues which need to be addressed, in parallel". This leads to triaging issues and figuring out which issues need to be addressed in what order. Even assuming perfect efficiency, the overhead of this sort of project management in-and-of-itself requires an allocation of resources which would otherwise be spent addressing said issues. The netCDF team does not enjoy perfect

project management efficiency.

The reduced resources limit the amount of effort that can go into directly implementing new features in netCDF. Furthermore, time is spent between purely technical tasks and other, equally important but nebulous tasks such as community maintenance and support, project management, and research into emergent technologies and how they can be used to meet the needs of our community.

Roadmap

The following items are prioritized in the medium-to-long term:

- Documentation
    - NetCDF has added a lot of functionality that has not been documented as well as we would like.
    - The documentation organization isn't very good.
    - The NUG needs to be modernized.
        - This has been ongoing for quite a while
- More reliable Amazon AWS access via ncZarr + documentation.
- Integration of cloud-HDF functionality.
- Continued bug fixes, optimization.
- Evaluation of the next generation of emergent technologies; we have had great success in the past anticipating what will be useful to our community. We want to continue this success.

Information specific to netCDF-Java can be found in the THREDDS Status Report.

Prepared *March 2026*

Status Report: Information Technology

September 2025 - March 2026

Mike Schmidt & 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 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

- As of the beginning of 2026, Unidata has turned over Windows and MacOS operations to UCAR's Enterprise IT (EIT) group
- Unidata IT staff continue to update and patch server systems to optimize security, productivity, and uptime.
- 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 competed and/or anticipated are;
  - centralized Identity and Access Management (IAM, coming very soon)
  - centralized end-user security agent (Palo Alto Cortex XDR, ongoing updates)
  - updating hardware infrastructure as required

WIth the change in management of the NCAR Wyoming Supercomputing Center (NWSC) announced, and the potential changes impacting the Mesa Lab Data Center (MLDC), there is a great deal of uncertainty about both primary data centers utilized by Unidata. Due to the volumes of data moved by the IDD, it seems unlikely this activity could be affordably relocated to the commercial cloud.

Local and cloud backups have been implemented redundantly on all Unidata end-user devices. Unidata IT can remotely manage Unidata client devices when they are connected to the Internet through UCAR's VPN service.

Daily, we continue efforts to keep services and systems secure which takes consistent attention and occasional herculean immediate 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 *Match 2026*

Status Report: Support

September 2025 - March 2026

Jennifer Oxelson, UPC Staff

Executive Summary

- We are working to address deficiencies in our software documentation to make the information they contain more user-friendly and discoverable.

Questions for Immediate Committee Feedback

What methods for seeking support or engaging in discussions within a community have you used and found particularly effective? E.g., Discourse? Forums? GitHub Issues or Discussions?

Activities Since the Last Status Report

Improving online self-support availability

1. **Software Documentation Improvement Underway**
   We are working to address deficiencies in our software documentation to make the information they contain more user-friendly and discoverable.

Note: compiling of support metrics for staff is on hold while we conduct a comprehensive review of our metrics gathering process.

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
