

Fall 2026 Community Advisory Committee Meeting

(Times are Mountain Standard Time)

Wednesday, September 9, 2026 9:00am – 4:30pm MT

[Meet online and in NCAR ML-239 (Mesa Lab – Damon Conference Room)]

9:00 – 9:30 Welcome & Introductions (Chairs/ Ryan May)

- Setting the Stage for the New Committee (Mohan)
- New Member Welcome
- Transitioning Member Appreciation (Certificates)
- Round the Room Introductions (Name, Institution, Highlights)

9:30 – 10:00 Director's Opener (Mohan Ramamurthy)

10:00 – 11:00 Staff Portfolio Lightning Highlights (UPC)

11:00 – 11:15 Interactive Connection Poll and Break

11:15 – 12:00 Portfolio Lightning Highlights Discussion (all)

12:00 – 13:10 Lunch

13:10 – 13:30 NCEP Update (Margaret Curtis)

13:30 – 13:50 NWS Update (Scott Jacobs)

13:50 – 14:00 Break

14:00 – 15:00 SESSION: NSF Unidata and AI (Co-led UPC & CAC)

- How are members leveraging AI in their work?
- How are members preparing students for AI workforce development?
- How could NSF Unidata maximize community benefit?
- Summary (output)

15:00 – 15:15 Day 1 Wrap-Up

15:15 – 15:30 Transition for DeSouza Award Presentation

15:30 – 16:30 DeSouza Award Presentation & Q&A with Dr. Eric Bruning, Mesa Lab Main Seminar Room (ML-132)

16:30 Adjourn

17:30 Working Dinner at [The Lowdown](#) ([map](#))

Thursday, September 10, 2026 9:00am – 4:30pm MT

[Meet online and in UCAR Mesa Lab – Damon Conference Room (ML-239)]

9:00 – 9:15 Welcome (Chairs/ Ryan May)

9:15 – 10:15 SESSION: Committee Model Implementation (CAC)

- CAC & Subcommittees Solicitation Process (output)
- Subcommittees Member Roster (output)

10:15 – 10:30 Break

10:30 – 10:50 Agency Update (TBD)

10:50 – 11:10 Agency Update (TBD)

11:10 – 11:45 Follow-Up “Future Directions for Earth Observations & Data Stewardship”
– What does NSF Unidata Need to do to Think Big in This Space?

11:45 – 13:00 Lunch

13:00 – 13:45 Business Development Update & Discussion (Ryan / CAC)

13:45 – 14:30 Day 2 Carry-Over Items, Next Steps, & Closing Remarks

14:30 – 16:30 Connection Activity (all)

16:30 Meeting Adjourned

Friday, September 11, 2026 9:00am - 1:00pm MT

__Strategic Advisory Subcommittee members only__

[Meet online and in UCAR FL4-1201 (Unidata Conference Room)]

9:00 - 9:15 Welcome (Chairs/ Ryan)

9:15 - 9:35 Budget Report (Richard Michel)

9:35 - 10:40 SESSION: Preserving Core Mission Amid Program Growth

- Grounding Principles & Guardrails that Keep Core Mission Intact as New Activities Introduced (Shaped by SAS)
- Report Out, Discussion
- Summary List of Principles & Guardrails (output)

10:40 - 10:55 Break

10:55 - 12:00 SESSION: Strategic Planning Subcommittee: Standing Up & Broadening the Plan

- Changes in the Landscape & Program
- Broadening the Plan Beyond the Core Proposal
- Summary List of Key Landscape Changes & Opportunities to Broaden the Unidata Strategic Plan that will be Provided to Subcommittee (output)
- Strategic Planning Subcommittee Member Roster (output)

12:00 - 12:30 Charter Updates & Considerations (SAS)

12:30 - 13:00 Sack Lunch & Closing Remarks (Chairs / Ryan)

Status Report: Community Advisory Committee Actions

September 2025 - March 2026

Unidata Program Center Staff

Actions from the Previous Meeting (March 2026)

Action 1

Implementation of New Committee Structure taken [CAC leadership, UPC]

Result

In Progress

Action 2

AI/ML if further feedback or thoughts on documents and prompts, ie gaps, etc. for community needs

Result

Further follow-up during Fall meeting

Action 3

Start to think about Fall meeting times [CAC, UPC]

Result

Complete

Action 4

Change in how data shared amongst countries and impact on Unidata and academic institutions

Result

Unknown

Status Report: IDD and LDM

April 2026 - September 2026

Stonie Cooper, Jennifer Oxelson Ganter, 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: Preparing for end of SBN, integrating other data gathering paradigms:
 - Developing directory pickup facility to replace the noaaportIngester for providing ease of crowd-source IDD data insertion.
 - Monitoring changes at NOAA wrt what will replace the SBN, and how the IDD will operate in that environment.
 - Continuing to mature the data cataloging mechanism.
- LDM: Continuing with source code updates and new feature additions:
 - Matching new root-less version with documentation and instructions for creating self-extracting installs without the need of the OS package manager.
 - New inotify-style LDM application for inserting data onto the queue being developed in response to potential cloud or crowd-sourcing of data.
 - Continuing to develop training materials

Dependencies, challenges, problems, and risks include:

- The lack of communication on the future of the SBN is a major obstacle to planning for the future of the IDD.

- 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.
- Continue to develop 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 *August 2026*

Status Report: THREDDS

March 2026 - September 2026

Sean Arms, Jennifer Oxelson, Ryan May, Ethan Davis

Executive Summary

The THREDDS project currently focuses on long-term sustainability while maintaining critical scientific data services. The releases of netCDF-Java v5.10.0 and TDS v5.9 delivered performance enhancements, bug fixes, and expanded support for object storage and GRIB feature collections. Metrics indicate steady adoption, with 5,330 unique TDS startups recorded between February and July 2026. We remain dedicated to encouraging migration from the unmaintained 4.x series to 5.x to bolster security. Furthermore, initiatives are underway to define the project's future direction, renew technical documentation, and integrate new cloud-native datasets into Unidata-maintained TDS instances.

Questions for Immediate Committee Feedback

Please see the very important note below regarding the future of the THREDDS project

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 findings from the NSF Unidata Community Assessment survey.
- Submitted an AMS Abstract To the 43rd Environmental Information Processing Technologies: FAIR and Open Data and Software within the Atmospheric and Ocean Sciences to Support Transparent, Reusable, and Efficient Research and Operations: "Evolving the Thematic Real-time Environmental Distributed Data Services (THREDDS) Project: Crossing the Language Barrier." We hope to see you all in January!

NetCDF-Java

- netCDF-Java v5.10.0 was released June 24th, 2026 See also: <https://github.com/Unidata/netcdf-java/releases/tag/v5.10.0>. Highlights include:
 - support for blosc compression
 - GRIB Feature Collection support for object storage
 - Software Bill of Materials (SBOMs) for uber jars (i.e. netcdfAll, ncldv, toolsUI) using the CycloneDX standard (driven by the OWASP Foundation)
 - new GRIB PDS template support
 - updated GRIB tables from the WMO
 - several bug fixes and updates to 3rd party libraries

TDS

- Version 5.9 was released on July 13th, 2026. This contained a number of bug fixes, security updates, and improvements. See also: <https://github.com/Unidata/tds/releases/tag/v5.9>. Highlights include:
 - Object store based GRIB Feature Collections are now supported in the TDS
 - does not yet include a turn-key process for generating GRIB index files
 - these could be generated prior to uploading GRIB files to the cloud
 - Blosc compression support the TDS (non-java code but no native libraries to install as they are shipped with the application)
 - updated datasetScan functionality to optionally exclude empty directories
 - a configurable default response to NCSS GRID requests
 - several bug fixes and updates to 3rd party libraries

Future of the THREDDS Project

NetCDF-Java is currently built on JDK 8 to preserve broad compatibility, but the evolution of the core Java ecosystem including developments like [Project Panama](#), [Project Loom](#), [Project Babylon](#), and [Project Valhalla](#), offers major opportunities for performance gains and codebase reduction through modern Java features as Java gains compiled code numeric performance within the JVM itself. These developments are backed by financial and developer resources from the open source community at large, Oracle, the Eclipse Foundation, Red Hat, Microsoft, IBM, AWS, and others, and represent a major under-the-hood overhaul for the language and its capabilities. Continuing feature development while targeting Java 8 is not a responsible path forward for the library, as Java 8 will soon reach end of life, although not before Java 11 and 17 (both within the coming year). At the same time, resources within the Unidata Program Center are precious few. The project is now faced with whether to transition to support modern JDK versions and features, or begin the sunset process of netCDF-Java with only security updates. Given the wide use of netCDF-Java in the community, sunsetting netCDF-Java would serve a major blow to the open source scientific Java community, as projects like the IDV, Panoply, the netCDF Toolbox for Matlab, McIDAS-V, and BioDataFormats (to name just a few) would be left without a critical, currently non-replaceable data access library. This is even more so true for the likes of the THREDDS Data Server, ERDDAP, and RAMADDA, which provide FAIR access to data across outside of the Java ecosystems. This is not a call for a decision on direction today, it is important to begin considering, with perhaps a special topic of discussion at the Spring 2027 committee meeting (although immediate thoughts are certainly welcome before then!). What information would you as the committee like to have to help guide this decision?

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.unidata.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

We are eagerly awaiting the release of RRFS (more to come in October!) As a reminder, the 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.

While not available publicly at the time this report was written, it is anticipated that a 14 day rolling-window IceChunk Store of MRMS Reflectivity will be available for evaluation (hosted in JetStream2 but publicly accessible everywhere). This dataset will also be available to users via thredds-dev.unidata.ucar.edu via a Java - Python gCDM bridge. The IceChunk Store is fed in near-real time via the IDD, which makes it unique amongst other IceChunk stores currently available.

Development

- netCDF-Java
 - We're actively working to extend the ability for netCDF-Java to support cloud-based Forecast Model Run Collections in the TDS
 - While we continue to improve support for Zarr v2, we're beginning to look into consolidated metadata, as well as Zarr v3 support.
- TDS
 - Our survey of academic institutions running the TDS is complete. We intend to share these insights with the community on the Unidata blog this fall.
 - 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 THREDDDS work:

- We are actively leading a proposal targeting [NSF 26-512](#): Unlocking Dataset Value for AI-Enabled Scientific Discovery ([AI Datasets](#)), tentatively titled “*Earth System Sciences Data Workbench for Dataset Discoverability*” (due date: November 4 2026)

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 (and soon, MRMS Reflectivity product), what other datasets would you like to see us target for development?

Relevant Metrics

THREDDDS Startup Metrics

	2026-02 — 2026-07	2014-08 — 2026-07
TDS Startup (unique IP address count)	5,330	54,701


	Total Servers	Information page updated
Publicly Accessible ¹ TDS count	91	50


Over the past 5 months, **5,330** unique IPs started up the TDS (February 2026 through July 2026). Since we’ve started tracking these metrics (v4.5.3, August 26th, 2014), we’ve seen the TDS startup from **54,701** unique IP addresses. There are currently **91** publicly accessible TDSs running “in the wild”. Of the **91** publicly accessible servers, **50** 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 v5.8 is

¹ “Publicly accessible” means we could find a top-level THREDDDS 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.


the dominant specific version running in the wild (although at least 1 has migrated to v5.9, which is the latest and was released July 13th, 2026). We continue to see a shift from 4.x to 5.x, which is super important, as 4.x is no longer maintained and subject to multiple security issues due to out-of-date 3rd party dependencies. If your loved ones are running TDS 4.x, please seek help on their behalf.






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). That number has increased to 111 publicly accessible TDSs (as compared to the 50 identified by the startup logs). The 2025 breakdown of sites by high level entity type is as follows:


108 Known TDS Sites*



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

108 Known TDS Sites*





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



Status Report: Science Gateway and Cloud Computing Activities

March 2026 - August 2026

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

Executive Summary

- Customized JupyterHub environments supported 227 students at nine institutions and workshops during the reporting period, extending a program that has served nearly 2,600 users since 2020.
- Unidata submitted its 2026–27 NSF ACCESS allocation request to sustain the CPU and GPU resources required by Science Gateway services on Jetstream2.
- Unidata is experimenting with an MRMS workflow using Icechunk, Zarr, Jetstream2 object storage, and THREDDS.
- Jetstream2 continued to host operational services, including AWIPS EDEX, THREDDS/NEXRAD, LDM/IDD infrastructure, and science gateways supporting external research collaborations.
- Collaborating with Jetstream2 to improve reliability, and security of Science Gateway infrastructure.

Questions for Immediate Committee Feedback

None at this time

Activities Since the Last Status Report

Jetstream2 2026-27 NSF ACCESS Grant Application

NSF Unidata has submitted a request to the NSF ACCESS program for 5,376,540 CPU SUs and 588,672 GPU SUs to support the upcoming Jetstream2 cycle, beginning October 1, 2026 through September 30, 2027. The amount of resources requested are approximately the same as the previous grant. If approved, this allocation will provide continued access to services including AWIPS EDEX, JupyterHub, THREDDS, and LDM/IDD nodes and ensure uninterrupted operation of CPU and GPU virtual machines supporting Unidata operations on Jetstream2.

MRMS JS2 Object Store

Unidata staff developed an experimental workflow for publishing Multi-Radar/Multi-Sensor (MRMS) reflectivity analysis data as cloud-native Zarr datasets backed by Icechunk and hosted on the Jetstream2 object store. The workflow ingests MRMS data produced by NOAA, uses LDM to transfer and transform the data, and stores the resulting datasets on the Jetstream2 object store. The data are then served through the standard interfaces of the TDS allowing users to leverage existing TDS workflows to access these large scientific archives. We plan to

transition this workflow into production during the next allocation cycle.

Sunsetting of RAMADDA on Science Gateway and Archival of ramadda-docker Project

Due to lack of interest for this project and to shift resources to other areas, the RAMADDA instance on the Science Gateway has been shut down. In addition, the ramadda-docker github project has been archived and all images associated with this project have been deleted on DockerHub. Going forward, we will use the RAMADDA deployment hosted at Unidata for situations that require the functionality offered by this server.

JupyterHub Activities

JupyterHub Servers for Workshops, and Spring Semester

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 spring semester of 2026 encompassing the length of this status report, 227 students at 9 academic institutions and various workshops have used NSF Unidata JupyterHub servers running on Jetstream2.

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.

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.

Science Gateway Repository Migration and 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 Jetstream2 (e.g., THREDDS, AWIPS EDEX). Over time, this repository has become unwieldy encompassing too many disparate concerns.

As part of a multiyear ongoing effort, we are 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.

Jetstream2 Allocation Management and Collaboration

NSF Unidata staff continue to manage our Jetstream2 allocation and collaborate with the Jetstream2 team. We 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 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:** A Magnum cluster becomes stuck during initial deployment requiring the cluster creation process to be restarted in an otherwise

- nearly completely automated Nushell driven workflow.
- **Outdated Magnum images:** Outdated cluster template images that can be more than a year old that may pose a security risk due to unpatched vulnerabilities.
- **Load balancer failures:** In some cases, load balancers have become unresponsive, rendering the cluster inoperable..

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, THREDDS (TDS), and the THREDDS Data Manager (TDM). In addition, we also maintain a security-hardened NSF Unidata Tomcat container inherited by 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.

More recently, we now maintain a Unidata Desktop container for our JupyterHub deployments so that users can access IDV and CAVE directly from the web browser and incorporate it into broader science gateway workflows alongside PyAOS notebooks.

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{Chastang2027a,
  title      = {{New Capabilities in the NSF Unidata Science Gateway}},
  author     = {Chastang, Julien and Espinoza, Ana},
  year       = 2027,
  month      = {Jan 10--14},
  booktitle  = {Proceedings, 43rd Conference on Environmental Information Processing
Technologies, 107th AMS Annual Meeting},
  address    = {Denver, Colorado, USA},
  note       = {Abstract submitted},
  organization = AMS
}
@conference{Espinoza2026,
  title      = {{How a Small Team Develops, Maintains, and Advances an Earth Systems
Science Gateway}},
  author     = {Espinoza, Ana and Chastang, Julien },
  year       = 2026,
  booktitle  = {Gateways 2026},
  address    = {Washington, D.C., USA},
  note       = {Abstract Submitted},
  organization = SGCI
}
```

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.
- Deploy large-language-model service on Jetstream2 with vLLM for IDV integration.

Relevant Metrics

Spring/Summer 2026 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.

<u>Spring 2026</u>		
AMS 2026 Student workshop	51	According to login spreadsheet
Seoul National University	33	Duseong Jo
Florida Institute of Technology	4	Milla Costa
SUNY Oswego	4	Scott Steiger
Millerville	34	
College of Dupage	28	
University of Georgia	23	
Florida State University	23	Christopher Holmes
Vermont State University	6	
Virginia Tech	0	
<u>Summer 2026</u>		
UCAR SOARS Internship	0	
UCAR PDWS	21	Drew

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, 2025 and September 30, 2026, NSF Unidata was granted an allocation on the NSF Jetstream2 (JS2) cloud valued at 5.1M+ CPU Service Units and 467k 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 August 16, 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	4.71M SU	399k SU	11.0 k SU/day	Over
GPU	467k SU	49.6k SU	417k SU	384 SU/day	Under

Based on these usage rates, we are over-budget on CPU SUs by ~130k SUs, 2.5% of our total allocated SUs. There are two main reasons for this:

- 1) Unexpected cluster deletion errors which needed JS2 administrative assistance left some machines idling away SUs
- 2) The amount of SUs requested by our 2025-2026 ACCESS allocation was based on estimates of auto-scaled node uptimes

In the past, Jetstream2 has been amenable to supplemental SU grants in the event more SUs are needed. We intend to apply for such a grant to ensure continued gateway uptime until our expected allocation renewal goes into effect on October 1, 2026.

Github Statistics*

Repository	Watches	Stars	Forks	Open Issues	Closed Issues	Open PRs	Closed PRs
science-gateway	5 (-1)	19 (-1)	13	6	168(+1)	9 (-4)	905 (+6)
tomcat-docker	9	67 (+1)	69 (-1)	0	43(+1)	1(-1)	100(+3)
thredds-docker	12	40	32(-1)	4 (+1)	127	1(+1)	195 (+7)
ldm-docker	8	15 (+2)	14(-1)	1	42	0	80 (+2)
tdm-docker	3	4	7	0	10	1(+1)	31

* Numbers in parentheses denote change from last stat report

Status Report: Data Standards and Technical Engagement

March 2026 - Sept 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 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. WIS2 is operational and rolling out globally. WIS2 is scheduled to fully replace the GTS by 2033. 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 v2 data models. As part of the GeoZarr discussions and taking advantage of the Zarr Convention framework, a NetCDF Zarr (NZ) convention has been proposed which would provide the mappings between netCDF and the Zarr v3 data models.

The Zarr v3 specification introduces a framework for extensions and one for conventions. While extensions provide a convenient way to add capabilities, they will add challenges for

implementers looking for interoperability.

Progress has been made on the following:

- Members of Unidata's netCDF team are coordinating with developers at DKRZ on their efforts to fully implement Zarr v3 support in the netCDF-C library.
- 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 as well as the GeoTIFF standard. GeoZarr will be developed within the OGC community standards process.

ESIP Activities

During the July 2026 ESIP meeting

- Attribute Convention for Data Discovery (ACDD) work continued on defining a governance and process for ACDD development
- Ethan Davis convened a session on WIS2 with others from NWS, UK MetOffice, and Environment Canada and helped convene another session on CF, netCDF, Zarr, GeoZarr.

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 2026 CF workshop ([website](#)) will be held 21-24 September 2026 at the ECMWF offices in Bonn, Germany and online.
- Ethan Davis continues serving as chair of the [CF Governance Panel](#).
- Members of the CF Governance Panel and CF Committees continue to participate in talks with the GeoZarr working group chairs.

OGC activities

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 *August 2026*

Status Report: netCDF

March - August 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 continue our engagement with the HDF Group, having monthly developer meetings with the HDF5 team, allowing us to stay ahead of the curve with any upcoming changes which might impact our community.
- Working, ongoing collaboration with developers at the German Climate Computing Center (DKRZ) to enhance [ncZarr](#). This collaboration has resulted in a lot of improvements for the ncZarr data format.
- 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 v4.10.1](#), 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, <https://hub.docker.com/r/unidata/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 *August 2026*

Status Report: AWIPS

April 2026 - Aug 2026

Tiffany Meyer

Executive Summary

At the time of writing this status update, I am currently working putting out a new release that should be out by the time the committee members meet. Version 23.4.-1 has been running 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. Version 23.4.3 is running on our back-up server and is the next release. This version of AWIPS is backwards compatible with previous version 23 builds and upgraded from Java 11 to 17. Additionally, this is the first release that has Mac installers for both intel and silicon.

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.

Unidata will be hosting a session at the AMS Student Conference where we will be giving an overview of Unidata products, have a panel where students can ask questions, and have some hands-on opportunities with different software. We will utilize the Virtual Desktop on a JupyterHub where students can connect to CAVE from their browser via the help of the Science Gateway group and Jetstream2. These JupyterHubs are also available in your classroom, just reach out.

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. 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.

New in v23.4.3-1 AWIPS:


- Upgrade to Java 17
- General FOSS version updates
- RRFS and REFS ready
 - New grib2 tables and grids
 - Support rotated lat/lon grid coverage
 - Updated menus
- AI-GEFS and GEFS model data now available
- Added Ensemble Tool
- Reconfigured logging to use Alertview
- Two installers for Mac - intel (x86_64) and silicon (arm64)
- Fixed Mac legend and colobar issue
- Fixed missing scales of radar bundles
- Fixed display of JASON Waveheight products
- Removed outdated satellite menus
- Removed obsolete lightning menus

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 20260101 - 20260821.

Daily AWIPS Users Connected to Public EDEX (Jul 1, 2025-Jun 30, 2026)



I typically have 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:

Service Quality



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
23.4.3-1	TBD	X	X	X	X	X	

As mentioned by the time the committee meets, a new build should be out. This includes EDEX (For RHEL8 systems), and CAVE installers available for Windows, Mac (intel and arm), Linux (RHEL/Rocky 8) and a virtual machine (VMWare). All source code is available. In addition, all python-awips notebooks were updated to work with the latest packages.

The naming convention the AWIPS Team follows for AWIPS releases tries to follow/reference the NWS. For example, version 23.4.3-1 is based off the NWS 23.4.3 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

3 - 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 is trying to stay up to date with the latest on CIRRUS/HIVE. 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

~2500 downloads (Mar 2025-Jul 2026)

Prepared *August 2026*

Status Report: IDV

March 2026 - September 2026

Yuan Ho

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 7.1 to be released on August 25, 2026. With this release, the IDV has successfully integrated the GEMINI api. We also add the Amazon cloud object store interface to browse and load the data directly. This upgrade improves performance 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 employs 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

__ IDV Object Store URL interface __

The IDV Data Choosers tab now includes an ObjectStore URL interface under General, so a cloud object can be specified the same way as an OpenDAP URL. The form lets you set the data source type, choose a host such as Amazon, optionally supply a profile and custom endpoint, then pick a bucket and object key. Pressing Add Source loads that object into IDV without downloading the archive first. This feature has been tested mostly in Amazon S3 bucket and can be extended to include Google, Microsoft, and others.




The screenshot shows a web form for configuring an Object Store URL. The form has the following fields and controls:

- Data Source Type:** A dropdown menu with the selected value "I'm Feeling Lucky".
- Host:** A dropdown menu with the selected value "Endpoint".
- Profile:** A text input field.
- Endpoint Host:** A text input field.
- Endpoint Segments:** A text input field.
- Bucket:** A dropdown menu with the selected value "noaa-goes19".
- Key:** A text input field containing the value "ABI-L1b-RadC/2026/160/17/OR_ABI-L1b-RadC-M6C02_G19_s20261601726185_e2026160172855".

Below the fields, there is a blue link that says "Press 'Add Source' to load the URL." and a button labeled "Add Source" with a help icon (question mark) next to it.

__IDV Amazon S3 Bucket Objectstore Browsing__

IDV can now browse Amazon S3 object stores directly from the Data Catalog Chooser, so users no longer need a local copy of the archive. Enter a bucket URL such as <https://noaa-mrms-pds.s3.amazonaws.com?delimiter=/> or <https://noaa-goes19.s3.amazonaws.com/?prefix=ABI-L2-CMIPF/&delimiter=/>, and IDV lists the bucket as a tree: S3 common prefixes appear as expandable directories and objects appear as files. Large listings are paged, folders can be shown newest-first or oldest-first, and a selected NetCDF or GRIB object can be opened in place through the cdms3: path, including use in NcML time aggregations. This puts public cloud holdings such as NOAA HRRR, NOAA Level2 Radar, NOAA MRMS, and GOES ABI on the same path as a local file, from browse to display.



__Gemini 2.5 AI process and IDV integrated displays__

IDV now includes a runGemini feature, shown as the Gemini Agentic Vision dialog, that sends a weather-analysis request to Google Gemini from inside the application. User must enter their own Gemini API key; IDV does not ship a shared key. A prepaid key can be created in Google AI Studio, then billed through Gemini API prepaid billing. After the key is stored, the user can attach archived weather maps and analysis files, choose a master prompt, reuse or clear prior instructions, and type an Instruction for Gemini Agent plus optional Knowledge Context. A Reset Gemini Key control replaces the stored credential. After OK, those inputs and images are sent with the user's key so the model can interpret the maps and return an analysis

without leaving IDV.

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 Jython formulas: Local Max/Min Identification

IDV now includes Jython functions that locate local maxima and minima on a grid and return them as point observations. `findLocalMaxMinAsPointObs2D(grid, radius, function)` scans a 2-D field, and `findLocalMaxMinAsPointObs3D(grid, level, radius, function)` does the same on a selected level of a 3-D field. Each grid point is compared with neighbors inside the given radius so that only true local highs or lows are kept, which is the usual way to mark H and L centers on surface pressure. The result is a set of points that can be plotted on the map, so closed-contour highs and lows can be labeled automatically instead of being picked by eye.

IDV Community Support

Naval Postgraduate School IDV weekly remote meeting

We provide remote support sessions as needed for 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.

Naval Postgraduate School IDV Seminar

We are offering one remote seminar on April 21, 2026 for a class of students at the Naval Postgraduate School. The seminar is an overview of IDV and new IDV features of integration of Gemini with IDV.

MSU Denver IDV Seminar

We are offering one remote IDV seminar on Jun 8, 2026 for a group of faculty member at the [MSU Denver](#). The seminar focused on the IDV Gemini integration discussion.

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

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 *August 2026*

Status Report: Python

March 2026 - September 2026

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

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. 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. Siphon development is low priority, but is supported by stable infrastructure and a simmering community engagement. MetPy development continues and 1.8.0 is slated for release late Fall 2026, featuring our C++-based performance enhancements, based on new build, test, and packaging infrastructure. MetPy's impact on science continues to grow, with 553 theses and peer-reviewed publications mentioning or citing MetPy, including 92 so far in 2026.

Questions for Immediate Committee Feedback

Nothing at this time.

Activities Since the Last Status Report

Python Training

Our present funded participation in Project Pythia is coming to a close this year. In June, we successfully led the 2026 Project Pythia Hackathon with participation from over 50 contributors, group leads, and facilitators. This is the last full-size in-person Hackathon event, though some leftover funds may enable one small event in the near future. The MetPy Cookbook (<https://projectpythia.org/metpy-cookbook/>) and all other Project Pythia resources will remain live online.

We are co-leading the delivery of a new Unidata session at the 2026 American Meteorological Society (AMS) Student Conference, composed of multiple panels and demonstrations. See the AWIPS and Community Services status reports for more information.

Our *Elements* track proposal to the last NSF CSSI solicitation to create Jupyter dashboards for classroom programming education was not selected for funding. Our sincere thanks for the strong collaborations from our committee and community, with special thanks to our co-submitter Dr. Casey Davenport at UNC Charlotte. We hope to retain the ideas and interest generated for future development or funding opportunities.

We supported the migration of Unidata's Machine Learning Foundations eLearning modules to MetEd, including maintenance of the course's GitHub resources.

MetPy

Development overall is a bit sluggish as priorities elsewhere in the Unidata program keep key resources occupied. Much of the focus remains on finishing and delivering the performance and scaling work for the last CSSI award, 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

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. Specifically, we have outstanding work to ensure compatibility with current releases of Xarray and Matplotlib.

Progress has been made on the following:

- MetPy 1.8.0 targeting late Fall 2026 release
- 553 total theses or peer-reviewed publications citing or mentioning MetPy, with 92 so far in 2026.

Siphon

Siphon 0.11.0 was released in July 2026, with the only user visible change to fix the Wyoming Upper Air client. Wyoming's original upper air server experienced a catastrophic failure, and a new server was brought online featuring new software infrastructure and a slightly changed API, which broke Siphon. Outside of this, Siphon has seen minimum updates outside of changes to keep CI infrastructure working smoothly. 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
- 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:

- Plan event for the Student Conference at the 2027 AMS Annual Meeting

- Release MetPy 1.8.0

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

- 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, 850 repositories and 119 packages depend on MetPy
 - 92 citations/mentions in 2026, 553 total
- Siphon
 - According to GitHub, 287 repositories and 36 packages depend on Siphon

Prepared August 2026

Status Report: AI & ML

March - August 2026

UPC Staff

Executive Summary

Unidata developers have advanced several key AI/ML initiatives, including the successful release of the *runGemini* feature in IDV 7.1, which is being followed up with the development of a context-aware IDV AI Assistant. Documentation for LDM and UDUNITS-2 was enhanced using AI to identify user challenges from support archives, while education and outreach efforts continued through the CU emergent ML group and a new article series on AI in Earth Systems science. We are also continuing to develop workflows that employ AI to streamline development activities like PR evaluation. Furthermore, the team is actively leading a proposal for the NSF “Unlocking Dataset Value” solicitation to improve scientific data discoverability.

Questions for Immediate Committee Feedback

- None at this time (in this report)

Activities Since the Last Status Report

Emergent AI/ML Activities

Ward has continued his work organizing an emergent ML technologies discussion group at CU, discussing AI in the classroom from both an educator, student, and staff perspective; this group is now in its third year. Ward has presented a different approach, based on several core principles, for using ML technologies in effective, reproducible ways that do not offload critical thinking or remove the human from the loop. He will be presenting a variation of this talk for the NAIRpilot Youtube channel, organized by CU/Rocky Mountain High Performance Computing consortium, and would love to present this information to anybody who will listen. Ward is collaborating with Sean Arms on a proposal for the ‘Unlocking Dataset Value for AI-Enabled Scientific Discovery (AI Datasets)’ solicitation.

AI In Education

2026 Summer Intern Tara Staresinic has [published the first article](#) in a series exploring the role of *AI in Education*, with a focus on its implications for the Earth Systems science education community. The remaining articles, which will examine different perspectives on AI and offer practical guidance for educators, are forthcoming.

Gemini and IDV Integration

In the 7.1 release, IDV now includes a *runGemini* feature that allows users to send weather-analysis requests to Google Gemini directly from within IDV. We collaborated with Professor Marty Baxter to develop a more effective master prompt for meteorological

analysis. Users can attach archived weather maps and analysis files, select a master prompt, reuse or clear previous instructions, and provide an instruction for the Gemini agent along with optional knowledge context. Gemini can then interpret the weather maps and return an analysis without the user having to leave IDV.

We are working on the IDV AI Assistant. It uses Retrieval-Augmented Generation (RAG) to support the design of an IDV AI Assistant. The goal is to combine the capabilities of a large language model with IDV's existing documentation, user guides, data descriptions, and domain-specific knowledge. This allows the AI Assistant to provide more accurate and context-aware guidance, help users identify appropriate data and visualization methods, explain IDV functions and workflows, and assist with weather-data analysis. By connecting the AI with IDV-specific knowledge and capabilities, it will provide an interactive assistant that can help users work with IDV more effectively while reducing the need to search through documentation manually.

AI-Assisted Documentation

AI has been leveraged to develop user-focused documentation for the LDM and UDUNITS-2 packages. AI agents reviewed public support archives, topical mailing-list archives, and GitHub issues to identify and catalog common installation and usage challenges experienced by users. The resulting documentation addresses gaps in the existing materials and will be made available shortly.

Proposal Development

We are actively leading a proposal targeting [NSF 26-512](#): Unlocking Dataset Value for AI-Enabled Scientific Discovery ([AI Datasets](#)), tentatively titled "*Earth System Sciences Data Workbench for Dataset Discoverability*" (due date: November 4 2026)

Progress has been made on the following:

- Shipping a runGemini feature in IDV 7.1 release
- Developing reproducible, effective workflows to leverage AI while keeping a human using critical thinking in the loop

Dependencies, challenges, problems, and risks include:

- Fitting in new activities alongside existing development is a challenge within the existing resources.

Ongoing Activities

We plan to continue the following activities:

- Developing AI integration within the IDV
- Exploring the use of AI to enhance documentation and other resources around projects
- Exploring uses of AI that improve development workflows

New Activities

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

- Incorporating an AI agent into the user support system has been identified as a low-effort, high-value opportunity. We plan to focus efforts in this area later this year.
- Submit a proposal to NSF's "Unlocking Dataset Value for AI-Enabled Scientific Discovery" solicitation

Prepared *August 2026*

Status Report: Support

September 2025 - August 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.

Incorporating an AI agent into the user support system has been identified as a low-effort, high-value opportunity. We plan to focus efforts in this area later this year.

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?

Ongoing Activities

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.
- 2. Support AI Agent**
Incorporating an AI agent into the user support system has been identified as a low-effort, high-value opportunity. We plan to focus efforts in this area later this year.
- 3. IDV Naval Postgraduate School IDV weekly support meeting**
The ID team provides remote support sessions as needed for the Lab Manager of the Interactive Digital Environmental Analysis (IDEA) Lab. More information is available in the [IDV status report](#).

Training

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

Relevant Metrics

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

Strategic Focus Areas

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

Prepared *August 2026*

Status Report: Information Technology

March 2026 - August 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

The Systems Group maintained and improved Unidata's technical infrastructure while advancing several strategic and development efforts. Key areas of work included:

Infrastructure, Operations, and Security

The Systems Group has been maintaining and supporting Unidata's core computing infrastructure, with an ongoing focus on system stability, lifecycle management, upgrades, patching, and security. This includes:

- managing Linux and other system updates,
- addressing vulnerabilities,
- monitoring system and network behavior,
- managing certificates and hardware,
- supporting infrastructure moves and decommissioning,
- and maintaining equipment and property inventories

Mike has also been addressing broader security and compliance requirements, including CMMC, and working to protect Unidata and JetStream resources from emerging threats and inappropriate use.

Web and Internal Services

Jen has been maintaining and modernizing Unidata's websites, web-based data services, and internal applications and services. Work has included application and platform upgrades, custom-module maintenance, configuration changes, troubleshooting, security improvements, and managing dependencies between applications and the underlying infrastructure.

Jen has also been evaluating the incorporation of AI agents into Unidata's support infrastructure.

Networking and Data Infrastructure

Mike has been maintaining and developing the networking and infrastructure that support Unidata's data services and external connections. This includes NOAAport, multicast networking, satellite-feed infrastructure, and other operational data connections, as well as planning for infrastructure changes and moves.

Mike and Stone have been evaluating IPv6 requirements and implementation, including associated technical and cost considerations, and working with UCAR and external partners on infrastructure issues that affect Unidata's data services.

Operational Support and Troubleshooting

The Systems Group has been providing ongoing operational support for Unidata staff, systems, and services. This includes investigating and resolving infrastructure, website, application, networking, and data-service problems; responding to support and security issues; coordinating with EIT and other UCAR groups; and handling routine system configuration and maintenance. The group has been balancing this day-to-day operational support with longer-term infrastructure and service improvements.

Software, Documentation, and Development Infrastructure

Jen has been working with Unidata developers to revamp UDUNITS-2 and LDM documentation.

The Systems Group has also been collaborating with staff on the management and cleanup of GitHub repositories, and consideration of organizational changes affecting Unidata's software-development infrastructure.

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

Status Report: Community Services

March 2025 - August 2026

UPC Staff

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:

- 2026 Summer Intern Tara Stareshin jump-started [Unidata's Instagram presence](#) in late July. Tara's internship has been extended, during which she will coordinate and manage Unidata's social media presence across all of its channels.
- Posts to the [News@Unidata](#) blog appear regularly, but not on a specific schedule. Some highlights:
 - [Next UDUNITS-2 Community Call - September 14](#)
 - [MetPy Mondays #336 - Accessing GOES Archives from S3](#)
 - [AI in Education Series - Introduction](#)
 - [MetPy Mondays #335 - Accessing and Plotting NEXRAD Level 2 Data](#)
 - [NSF Unidata Program Center is hiring a Community Services & Partnerships Manager!](#)
 - [NSF Unidata Joins Instagram](#)
 - [UDUNITS-2 Community Call](#)
 - [TDS version 5.9 released](#)
 - [MetPy Mondays #334 - Accessing and Plotting NEXRAD Level 3 Data](#)

- [MetPy Mondays #333 - How to Detect Highs and Lows](#)
- [netCDF-Java version 5.10.0 released](#)
- [MetPy Mondays #332 - Getting GFS Case Study Data from NCEI](#)
- [Reimagining NSF Unidata's Advisory Committees](#)
- [NSF Unidata eLearning Resources Now Available on MetEd](#)
- [MetPy Mondays #331 - Making an Emagram Plot](#)
- [Welcome Summer Intern Tara Staresinic](#)
- [Welcome Summer Intern Sarah Ravellette](#)
- [MetPy Mondays #330 - Making a Stüve Plot](#)
- [MetPy Mondays #329 - What's New in MetPy?](#)
- [Unidata's TDS Keeps Humans in the Loop for Metadata Quality Assurance](#)
- [Update on Community Access to NSF Unidata McIDAS-X](#)
- [NetCDF 4.10.0](#)
- [Register Now for the 2026 Pythia Hackathon](#)
- Community meetings and other announcements
- Updates to NSF Unidata's social media channels (Facebook, Bluesky, LinkedIn, Twitter/X, and now [Instagram](#)!) 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

With the retirement of Doug Dirks, Jennifer Oxelson Ganter has assumed responsibility for administering the Unidata website and coordinating blog content. Unidata staff members Once the new Community Services & Partnerships Manager is hired, responsibility for blog content creation and coordination will be evaluated.

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

- Proposed a new Unidata-themed event at the 2026 AMS Student Conference composed of multiple separate panel and demo sessions (Tiffany Meyer + Drew Camron, April) (accepted, development in-progress)

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

In no particular order:

NCAR Community Software Facility (CSF) Architecture Workshop	Drew Camron
Project Pythia	Drew Camron

UCAR Web Technical Governance Advisory Group (WeTAG)	Jennifer Oxelson Ganter
Mentoring in UCAR Mentorship Program	Jennifer Oxelson Ganter Sean Arms
UCAR Enterprise Change Control Board (ECCB)	Mike Schmidt
UCAR Computer Security Advisory Committee (CSAC)	Mike Schmidt
UCAR Network Coordination Advisory Board (NCAB)	Mike Schmidt
NCAR/CSU LROSE (LIDAR RADAR Open Software Suite)	Ana Espinoza Julien Chastang
JetStream 2	Ana Espinoza Julien Chastang Mohan Ramamurthy
2i2c Network: Open infrastructure for research and education	Ana Espinoza Julien Chastang
NASA Goddard Institute for Space Studies (Panoply)	Sean Arms
UDUNITS-2 Reboot Community	Ward Fisher Sean Arms Ethan Davis Jennifer Oxelson Ganter
HDF Group	Ward Fisher Sean Arms
OPeNDAP Group	Sean Arms Ethan Davis Ward Fisher
NCAR GDEX	Sean Arms
UCAR Software Engineering Assembly (SEA)	Ward Fisher
CU AL/ML Technologies in the classroom across disciplines	Ward Fisher
AI/ML Reading Group	Ward Fisher
Zarr Developers Committee	Ward Fisher
UCAR LCPO Shared Expertise	Ward Fisher
German Climate Computing Center (DKRZ)	Ward Fisher
AMS EIPT (Environmental Information Processing Technologies) Committee	Tiffany Meyer

SSEC McIDAS-V Group	Yuan Ho
Naval Postgraduate School Interactive Digital Environmental Analysis (IDEA) Lab	Yuan Ho
MSU Multi-agent Integrated Tutoring System	Yuan Ho
Midwest RISA Co-learn on Indigenous Topics	Stonie Cooper
UIUC/I-GUIDE	Mohan Ramamurthy Inken Purvis
COMET LMS	Tanya Vance Mohan Ramamurthy
UCAR/UCP Business Development	Ryan May Mohan Ramamurthy
UCAR Friends of the National Center	Ryan May Mohan Ramamurthy
AMS GPP Committee	Mohan Ramamurthy
World Meteorological Organization (WMO) WIS2	Ethan Davis Stonie Cooper
Open Geospatial Consortium (OGC)	Ethan Davis
Earth System Information Partners (ESIP)	Ethan Davis
CF Conventions	Ethan Davis
NCZarr/Zarr Specification Efforts	Ethan Davis Ward Fisher Sean Arms
Earthmover PBC	Ryan May Ward Fisher Ethan Davis Sean Arms

Learning Services

Reimagined Science Gateway / Educational Hub

- For Science Gateway and Integrated Educational Hub information, please see the [Cloud Computing Activities](#) report for more details.
- Unidata eLearning content has been migrated to COMET's learning management system (LMS), [MetEd](#).

Project Pythia Hackathon

- The Project Pythia Hackathon took place on June 15-18, 2026 in Boulder, CO at the NCAR Mesa Lab. For more information, please see the [Python Status Report](#).

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 year months, we plan to organize or take part in the following:

- A **Naval Postgraduate School IDV Seminar** will be held on April 21, 2026. More information is available in the [IDV status report](#).
- A **MSU Denver IDV Seminar** will be held on June 8, 2026. More information is available in the [IDV status report](#).

Relevant Metrics

Social media statistics, August 20, 2026

2026 Summer Intern Tara Staresinic jump-started [Unidata's Instagram](#) presence in late July!

Social Media Channel	# of followers/subscribers	Notes
Instagram	83	Brand new since late July 2026!
Bluesky	504	(# of followers up from 486)
Twitter/X	1975	(# of followers down from 2003 -- We just recently resumed reposting to X)
Facebook	996	(# of followers up from 992)
YouTube	4542	(# of followers up from 4470)
LinkedIn	639	(# of followers up from 567)

Unidata eLearning statistics

NSF Unidata successfully transferred our eLearning content to [COMET's MedEd LMS](#). The new platform will provide enhanced tracking and reporting of engagement with our educational

resources and learner activity. Metrics reporting will resume once the new Community Services & Partnerships Manager joins the team.

NOTE: Other metrics reporting will resume once the new Community Services & Partnerships Manager joins the team.

Prepared *August 2026*