

Users Committee/Strategic Advisory Committee Joint Session

Agenda

Monday, October 16, 2017

Users Committee Session: Unidata Conference Room

09:00 - 10:00 New Members Orientation (Josh Young/Staff)

*** Returning Members Convene at 10AM ***

10:00 - 10:15 Break/Full Committee Convenes

10:15 - 10:45 Welcome and Introduction

10:45 - 11:15 JetStream Science Gateway (Julien/Rich Signell)

11:15 - 11:45 Proposed Triennial Themes (Kevin G., Victor G., and Warren P.)

11:45 - 1:00 Lunch

1:00 - 1:45 Around the Table Reports (Committee)

1:45 - 2:20 Status Reports (Committee/Status)

2:20 - 2:50 NOAA CONDUIT (Becky Cosgrove)

2:50 - 3:00 Break

3:00 - 3:30 GOES Data access

3:30 - 4:00 MetPy developments (Ryan May)

4:00 - 4:45 Blue Skies (Committee)

4:45 Adjourn

6:00 Collaborative discussion on the day's proceedings over dinner at the NCAR Mesa Lab Cafeteria

Tuesday, October 17, 2017

Joint Session: EOL Atrium, FL-1

09:00 - 10:15 Director's Report (Mohan Ramamurthy)

10:15 - 10:45 NASA Update (Chris Lynnes - remote)

10:45 - 11:00 Break

11:00 - 11:30 Cloud Activities

11:30 - Lunch

1:00 - 1:30 Committee Recruitment
1:30 - 2:30 NOAA Big Data Project (remote presentation from Ed Kearns)
2:30 - 2:40 Break
2:40 - 3:15 Unidata and Workforce Development, Education, and Training
3:15 - 3:25 NOAAPORT Delivery of GOES 16 ABI Imagery (Tom Yoksas)
3:25 - 3:45 Discussion of future workflows (GOES workflow - Victor Gensini)
3:45 - 4:00 Transition to Auditorium
4:00 - 5:00 DeSouza Award Presentation entitled *Visualizing Visualization*
by 2017 Award Recipient Kevin Tyle
5:15: UCP Directorate Reception for SAC and Users Committee

Wednesday, October 18, 2017

Strategic Advisory Committee Session: Unidata Conference Room

9:00 - 9:15 Welcome and Administrative Items
9:15 - 10:15 Budget Report (Terry Mitchell-Sur)
10:15 - 10:45 Strategic Plan Update (Mohan Ramamurthy)
10:45 - 11:00 Break
11:00 - 11:30 Funding (Mohan Ramamurthy)
11:30 - 12:00 Business Development (Josh Young)
12:00 - 12:15 Hand-off of Chair (Bill Gallus/Sepi Yalda)
12:15 - 1:00 Review Actions/Lunch
1:00 Adjourn

Status Report: AWIPS and GEMPAK

April 2017 - October 2017

Michael James

Software Releases

AWIPS 17.1.1-1 (July 2017)

- Support for **GOES-16(R)** provisional ABI imagery. New GOES East Full Disk map scale added to D2D <http://unidata.github.io/awips2/cave/goes-16-satellite/>
- Support for **GOES-16(R) GLM** (Geostationary Lightning Mapper) as it comes online - glm data are not currently being ingested but will be added to edex-cloud as Unidata begins to test distribution feeds internally.
- **WarnGEN** and **TextWS** have been refactored to work for any and all WFO sites, with CWA geometries pre-generated in the build. I've made available a document and screencast walkthrough for using WarnGen to create (but not send) warning products: <http://unidata.github.io/awips2/cave/warnngen/>
- The **NSHARP** plugin is now written entirely in Java, no longer dependent on legacy NAWIPS/GEMPAK libraries.
- The **Manual Ingest Endpoint** has been moved from `/awips2/edex/data/hdf5/manual` to `/awips2/data_store/ingest` to allow for manual file ingest on an nfs-mounted raw data store.
- **LDM** pattern actions added for **TPCWindProb**, **RTOFS**, **Sea Ice**, **CFS-DNG**, **NAEFS**.
- 3-way right-click toggle in CAVE to cycle the resource stack through displaying **Products**, **Maps**, and **Current Time**.

AWIPS 17.1.1-2 (July 2017)

- Fixes Open Bundle dialog error in CAVE.
- Preserves current map views by opening saved bundles in their own editors.
- **File > Delete Bundle** menu item to manage saved bundles.
- Cleanup **GOES-R Derived Motion Wind** menu for provisional products.
- **awips2-gfesuite** RPM replaces gfesuite-server and gfesuite-client to avoid package conflict on standalone systems.
- **macOS** bundled JRE jspawnhelper set to executable to enable CAVE PyDev config.
- **gridslice.so** added to **awips-python.pkg** to support NumPy array slicing in CAVE (previously prevented isentropic grid bundles from loading)
- The python dependency package is now called "**awips-python.pkg**" rather than "**awips-python-jep.pkg**" and include Jep, gridslice, and other required libraries for macOS.

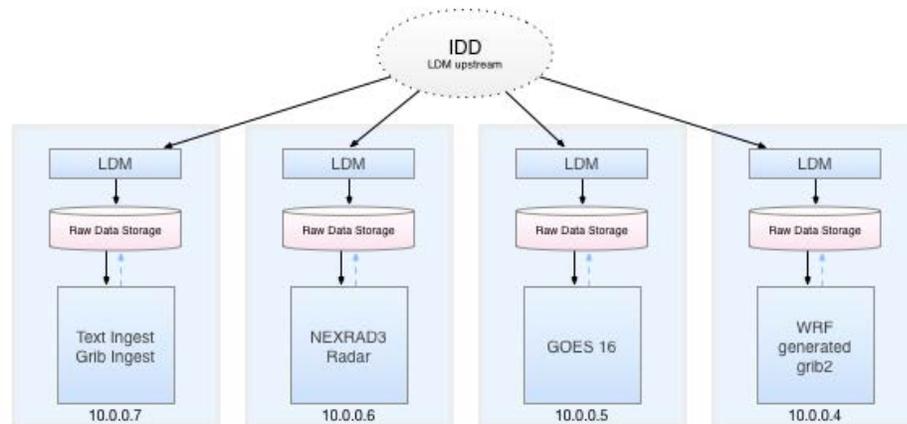
AWIPS 17.1.1-3 (August 2017)

- Updates to **GFE client** to enable grid editing in macOS, if client connects with default localization site OAX.
- Removed all **ISC fax/send** dialogs from the GFE perspective.
- Maps and NCEP database RPM fix to include shell login with su psql calls.
- Removed all remaining unused **Dat** (local data), **Radar Server** dependencies, **LDAD**, and **OHD** plugins which reduced the EDEX RPM build time considerably.

AWIPS in the Cloud

Unidata continues to provide access to real-time AWIPS data in the cloud thanks to the XSEDE Jetstream grant outlined in the Cloud Activities status report.

- edex-cloud is now distributed across 5 Jetstream server, with 4 decoding nodes store data to a single database and request server instance.
 - ingest/ingestGrib - default IDS|DDPLUS|FNEXRAD|NGRID|CONDUIT|NIMAGE|UNIWISC
 - ingestGoesr - NIMAGE
 - ingestRadar - NEXRAD3
 - ingestExternal - NAM3km via NCEP ftp; experimental ingest of AWS public datasets; WRF?



- This configuration *separates decoding resources from request resources*, and avoids PyPIES and LDM both writing to disk simultaneously.
- Remote PyPIES (HDF5) access is allowed by opening ports 9581 and 9582 in iptables
- Remote PostgreSQL is allowed by opening port 5432 in iptables and specifying unique IP address in **pg_hba.conf**, and requiring clients to supply passwords (sent with md5 encryption) - outlined at <http://unidata.github.io/awips2/edex/distributed-computing/>
- This is a different distributed configuration from what was running in the Azure cloud in 2016, which used a shared drive to separate LDM ingest from EDEX decoding (not possible in Jetstream).
- "**edex password <newpassword>**" command to update **hibernate.cfg.xml** files (metadata, maps, etc.) - after manually changing password from default on the

database/request server (cmd line from psql) - not essential since pg connections from specific IPs are defined in pg_hba.conf and iptables, but good practice.

- **yum groupinstall awips2-ingest** and **yum groupinstall awips2-database** available (former does not include postgresql, psql, PyPIES/hdf5, maps-database, topo, etc;; latter does not include LDM, ingest, ingestGrib)
- **“edex start ingest”** and **“edex start database”** commands available
- What this means: EDEX lives up to its SOA name, totally modularized... we could spin up a VM every 6 hrs to run a WRF domain, decode output and persist it to edex-cloud before shutting down the VM.
- Easier to manage clean reinstalls of EDEX which don't remove and reinstall the database and its many large files and slow progress.
- With postgres SSL certificates in the pipeline (17.3.1) we could even allow trusted users to spin up edex decoder instances for decoding their own datasets, which are saved to edex-cloud and available in CAVE and Python.

OS X Support

Unidata has supported CAVE for OS X since 15.1.1. In last year's status report I outlined three obstacles which were being addressed to make the Mac client more functional, and all three have (mostly) been fulfilled

1. *Jep and PyDev configuration has not worked correctly in previous versions (15.1.1, 16.2.2). Fixing this will allow for derived parameter functions to be run from CAVE (which is running locally-installed Python) - **Complete:** Jep has been bundled in all 17.1.1 releases and is correctly rendering derived parameters and wind representations*
2. *GEMPAK and NSHARP libraries should be built without static gcc references, we should not expect users to have installed libgfortran via homebrew or macports in order to plot GEMPAK-like grids - **Partially Complete:** CAVE now uses pure Java for the NSHARP plugin, no longer dependent on GEMPAK/NAWIPS libraries. AWIPS GEMPAK-like grid rendering is still dependent on legacy GEMPAK/NAWIPS libraries.*
3. *CAVE for Mac Sierra (10.12) must now be signed for distribution outside the App Store, otherwise the application will be quarantined by Gatekeeper Path Randomization, which prevents applications delivered outside of the store from running by normal means (taskbar icon or finder) - **Complete:** Unidata CAVE for macOS is now signed with the UCAR developer certificate and able to install and run without being quarantined (though you will still be warned about its origin since it is downloaded outside of the App Store).*

Python Data Access Framework (python-awips)

The [standalone Python Data Access Framework](#) (DAF) for retrieving data from an AWIPS EDEX server has been updated with changes from Vlab 17.1.1 repositories, and a Python 3-compliant version has been made available by PyPI (installable via pip) and integrated into the Python Training Workshop.

However, irregularities have been noticed in what datasets and parameters are available via the Python DAF between EDEX versions 16.2.2 and 17.1.1. Further changes coming in 17.3.1

may resolve these issues but it's ongoing work to figure out how AWIPS Java and Python correctly communicate.

Material Design Documentation Framework

I maintain a Material Design documentation package called **mkdocs-unidata** (a fork of <https://github.com/squidfunk/mkdocs-material>), which takes markdown files, organized as partials and includes), and outputs a rich material design rendering on GitHub Pages.

- <https://mjames-upc.github.io/mkdocs-unidata/>
- pip install mkdocs-unidata
- Used for the Unidata AWIPS User Manual <https://unidata.github.io/awips2/>
- Soon to be used for a redesigned GEMPAK User Manual

GEMPAK/NAWIPS

GEMPAK 7.4.1 (September 2017)

<https://github.com/Unidata/gempak/releases/tag/7.4.1>

- Grib2 table updates for moisture, momentum, mass, radiation, hydrology, cloud/aerosol, more.
- Removed repetitive dcmetr log messages.
- Pattern actions in ldm/etc/templates/pqact.gempak_decoders_grid for various NamDNG grid and the National Blend of Models.
- Alias/mod_res updates, increased max grids in file for gfs, nam, National Blend, SST, URMA, others.
- Bug fix for AWIPS grid retrieval in scripts/python/GridDataRetriever.py

GEMPAK 7.4.0 (August 2017)

<https://github.com/Unidata/gempak/releases/tag/7.4.0>

- Increased nex2img maximum GINI size
- Maps and bounds file operational updates from NAWIPS 7.4.0
- Grib table and pqact updates for 2.5km GFSMOS
- Fixed G-AIRMET issues when more than 20 points are drawn below the centroid of a smear or snapshot polygon
- Bug fix in ensemble diagnostic function ENS_CVAL
- Added new GEMPAK grid diagnostic function: neighborhood functions, developed by SPC
- Updated HYSPLIT restore files for parameter name changes
- Fixed nmap2 and gdplot2 bug where program was unable to display global wrapped CED grids for subareas including poles
- Fixed prob2cat bug that occurs while composing convective outlook, and bug where lines are drawn outside international boundaries
- Fixed dccsig when handling isolated or lines of thunderstorms

- Linux binaries built with Docker containers on Travis CI for CentOS 6, CentOS 7, Fedora Core 26, Ubuntu 16.04.

GEMPAK 7.3.2 (June 2017)

<https://github.com/Unidata/gempak/releases/tag/7.3.2>

- Further gdradr support for dual-pol products
- dcgrib decoding fix for NOAAport NAM212.
- g2vcrdwm06.tbl, g2varswmo0.tbl update for dcgrib.
- gdplot fix for missing titles

Ongoing/Future Activities

- Dockerize gempak on top of LDM.
- Dockerize and deploy gempak container from travisci.
- Implement AWS API for public data set retrieval via EDEX (cron or quartz)
- Update qpid from 0.32 > 6.1.4.
- Update AWIPS postgres/database to use SSL certificates for remote connection.
- Create unique, deployable ingest modes for Radar, Goesr, Grid, other data sources, essentially mini-instances of EDEX for each datatype focus.
- Investigate Appliance for further Linux distro support for CAVE.
- AWIPS 17.1.1-4 (next release candidate)
 - Consolidate EDEX RPMs (see below)
 - Separate decode/ingest from database/request installations and configurations (see below)
 - Add NationalBlend model to LDM ingest and CAVE display.
 - GOES-16 Imager, Derived Motion Wind, and Derived Product CAVE menu additions.
 - Remove Hydroapps completely from EDEX.
 - Remove awips2repo netcdf packages which caused conflict with epel repositories.
- Continuous Integration
 - EDEX and Qpid RPM builds run on remote push by Travis CI, using Centos 6 Docker container to install AWIPS Development Environment (ADE) packages (Java, Eclipse, Python, Ant), clone the multiple awips2 repositories from github.com/Unidata, and build a subset of RPMs.
 - It's not feasible to automate these builds for CAVE given the manual PyDev configuration required within Eclipse on the desktop - need to determine if we can bootstrap packages a PyDEV config inside Eclipse for RPM distribution.
- RPM Package Consolidation
 - Last year's AWIPS release 16.2.2 consolidated all awips2-cave-<pluginname> RPMs into two packages: awips2-cave and awips2-cave-wrapper.
 - Current release candidate 17.1.1-4 consolidates all awips2-edex-<pluginname> RPMs into the existing awips2-edex-dataplugins package. Rather than installing

(awips2-edex-binlightning, awips2-edex-bufr, awips2-edex-grid, etc.), it's a single RPM: awips2-edex-dataplugins.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

- 1. Enable widespread, efficient access to geoscience data**
Both AWIPS and GEMPAK are freely available, and both incorporate LDM/IDD technology for accessing geoscience data. The cloud-based EDEX data server continues to see widespread use and growing adoption. More and more datasets continue to be added to the server as Unidata deploys more decode/ingest nodes.
- 2. Develop and provide open-source tools for effective use of geoscience data**
Both AWIPS and GEMPAK are open-source, and while GEMPAK is now in maintenance mode, AWIPS is continuously being developed.
- 3. Provide cyberinfrastructure leadership in data discovery, access, and use**
Unidata is the only known entity to provide a freely-available and non-operational version of the AWIPS software package. Unidata continues to find and make available new datasets through the AWIPS project.
- 4. Build, support, and advocate for the diverse geoscience community**
Using LDM/IDD technology to provide access to real-time meteorological data; providing visualization tools for data analysis.

Prepared October 2017

Status Report: *Cloud Computing Activities*

April - October 2017

Julien Chastang, Ward Fisher, Michael James, Ryan May, Jen Oxelson, Mohan Ramamurthy, Christian Ward-Garrison, Jeff Weber, Tom Yoksas

Activities Since the Last Status Report

Unidata Science Gateway on Jetstream

The screenshot shows a web browser window with the URL `science-gateway.unidata.ucar.edu`. The page has a dark blue header with the text "Table of Contents". The main content area is white and contains the following sections:

- 1 Introduction**

Welcome to the [Unidata Science Gateway](#) on the [XSEDE Jetstream Cloud](#). As part of [Unidata's 2018 Five-year plan \(PDF\)](#) "Unidata 2018: Transforming Geoscience through Innovative Data Service", Unidata is exploring the use of cloud computing. A collection of Unidata related technologies can be found here for our community to make use of directly or with client applications described further on. Gateway users, coupled with XSEDE HPC resources, can achieve complete end-to-end scientific computing workflows.
- 2 JupyterHub**

[JupyterHub server on Jetstream](#)

[JupyterHub](#) is a technology that can be used to serve programmatic, interactive notebooks to a class of students or for scientific researchers. An [experimental JupyterHub server](#) is running on Jetstream containing Unidata Jupyter notebook projects:

 - [Unidata Python Workshop](#)
 - [Unidata Notebook Gallery](#)
 - [Unidata Online Python Training](#)

This JupyterHub server is currently experimental. If you would like to be granted access, please contact support@unidata.ucar.edu.
- 3 THREDDS Data Server**

[TDS installation on Jetstream](#)

The Unidata [THREDDS Data server \(TDS\)](#) is a web server that provides metadata and data access for scientific datasets, using a variety of remote data access protocols. A [TDS](#) is available on Jetstream at <http://thredds-jetstream.unidata.ucar.edu/thredds/catalog.xml> supplying a good portion of the data available on the [IDD](#) with a five day archive.

This TDS can be accessed from the [IDV](#) or from Python with the [netCDF-Python](#) or [Siphon](#) APIs or any THREDDS client (e.g., PyCSW).
- 4 AWIPS EDEX**

EDEX installation on Jetstream: edex-cloud.unidata.ucar.edu

AWIPS is a meteorological display and analysis system used by the National Weather Service.

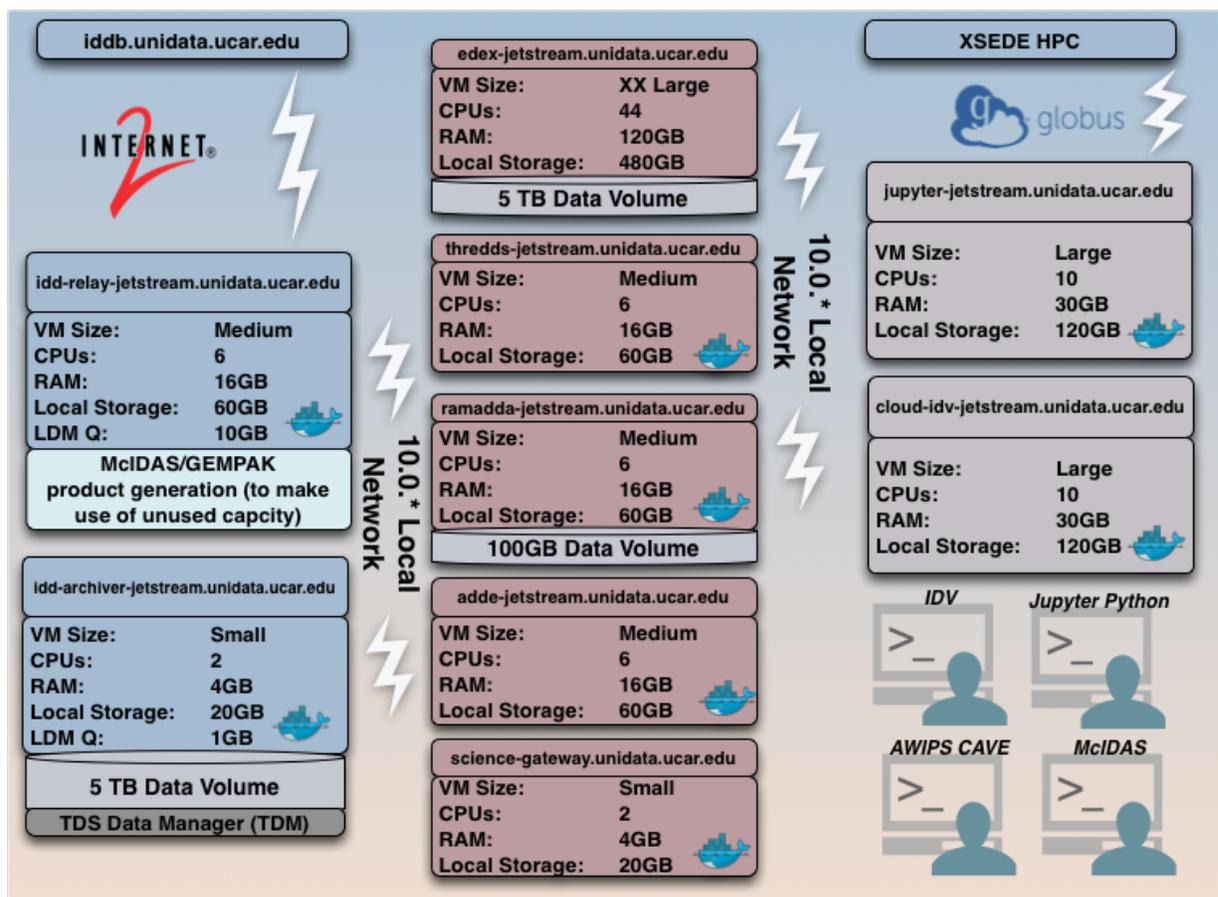
Date: 2017-09-27 16:00:39 MDT

Building upon our previous containerization efforts, we are developing a Unidata Science Gateway on NSF-funded XSEDE Jetstream Cloud: <http://science-gateway.unidata.ucar.edu/>. A

collection of Unidata related technologies can be found here for our community to make use of directly or with client applications such as the IDV. The following resources are available on this gateway:

- An experimental JupyterHub server containing Unidata Jupyter notebook projects
- A TDS supplying a good portion of the data available on the IDD with a five day archive.
- EDEX server
- RAMADDA containing IDV bundles that retrieve data from Jetstream data servers.
- Two LDM nodes
- ADDE Server
- IDV Jetstream plugin allows easy access to Jetstream installations of the TDS, RAMADDA and ADDE from the IDV.

Gateway users, coupled with XSEDE HPC resources, can achieve complete end-to-end scientific computing workflows. We presented this work at the [ESIP Summer Meeting 2017](#).



Dependencies, challenges, problems, and risks include:

- We aim to use cloud technologies such as Kubernetes and [Zero to JupyterHub](#) for JupyterHub so that we can scale for on-demand usage, in a classroom setting, for example. We will collaborate with our partners at XSEDE to arrive at this outcome.

- We will do more to promote and advertise this gateway.
- Unidata staff is working closely with Unidata system administrators to ensure cloud VMs (especially on Jetstream) are adhering to Unidata security standards.

EDEX in the Cloud

Unidata maintains an EDEX data server on Jetstream to ingest and serve real-time AWIPS data for rendering by the [CAVE client](#) and the [python-awips data access framework](#): `edex-cloud.unidata.ucar.edu`. This EDEX server has successfully supported several AWIPS workshops, and is used by CAVE clients in the Unidata community.

Nexus Server on Jetstream

Unidata is running a Nexus Server on Jetstream for the distribution of netCDF-Java artifacts (e.g., `netcdfAll.jar`, `toolsUI.jar`, `ncIdv.jar`): <https://artifacts.unidata.ucar.edu>. netCDF-Java documentation is also hosted at that location.

Transitioned from XSEDE "Startup" to "Research" Allocation on Jetstream

To further investigate how the Unidata community can benefit from Unidata technologies in the cloud, Unidata obtained a large XSEDE "Research" grant on the Jetstream cloud-computing platform worth \$425,000 in cloud computing resources. The Extreme Science and Engineering Discovery Environment (XSEDE) five-year, \$121-million award is a National Science Foundation supported project. In the last six months, we completely transitioned our research and development from our initial "Start Up" allocation to our "Research" allocation.

Docker Containerization Unidata Technology

We have been employing Docker container technology to streamline building, deploying, and running Unidata technology offerings in cloud-based environments. Specifically, we continue to refine and improve Docker images for the IDV, LDM, ADDE, RAMADDA, THREDDS, and Python with Unidata Technologies. We have been experimenting with these Docker containers in the NSF XSEDE Jetstream cloud.

Progress has been made on the following:

Containerization efforts are currently in maintenance mode with most of the initial development completed. The TDS docker container is the most active of our containerization efforts. We continue to receive open-source contributions. It is also our most popular container on DockerHub.

Dependencies, challenges, problems, and risks include:

It is unlikely that most of our community will use these containers directly. Rather they will be leveraged by experts on behalf of the community, or they will be abstracted from users by

being integrated into a user-friendly workflow. For example, on Jetstream we have a JupyterHub server currently in development: <https://jupyter-jetstream.unidata.ucar.edu>. This server was deployed with the aid of cloud computing technologies including Docker. These details, however, are hidden from the user.

In addition, there are overlapping (perhaps, competing or complementary) technologies such as Ansible that are emerging alongside Docker that need to be investigated.

2017 Modeling Research in the Cloud Workshop

Unidata obtained supplemental funding from NSF to host the [2017 Modeling Research in the Cloud Workshop](#), May 31 - June 2 2017. The purpose of the conference was to facilitate an in-depth discussion of the myriad aspects and formulate approaches for integrating cloud computing capabilities into the weather and climate prediction landscape and discuss the significance of such integration for advancing discoveries.

After the conference, we had a discussion with Kevin Tyle (University at Albany - SUNY and SAC member), Carlos Maltzahn (University of California, Santa Cruz, Big Weather Web Lead), and John Exby (formerly of NCAR, RAL) on Jetstream cloud use. We initiated a "Startup" grant on Jetstream for experimentation purposes. We also educated the group on the use of the TDS Docker container with a demonstration at Unidata. Subsequently, Kevin Tyle, and Julien Chastang collaborated on the use of the TDS Docker container.

Ongoing Activities

Amazon Web Service Activities and NOAA Big Data Project

NOAA Big Data Project

- In collaboration with Unidata, NOAA is delivering 20+ years of NEXRAD Level II data via Amazon Web Services. LDM and THREDDS Data Server (TDS) software are being employed to deliver these data.
- Started transferring GOES-16 data to the Amazon cloud S3 bucket.
- TDS on AWS for level II NEXRAD (For .edu access only): <http://thredds-aws.unidata.ucar.edu/thredds/catalog.html>
- AWS Explorer (Public access): <https://s3.amazonaws.com/noaa-nexrad-level2/index.html>
- Public Bucket for level II NEXRAD: <https://noaa-nexrad-level2.s3.amazonaws.com>
- Continue to populate the NEXRAD level II archive with real time data.
- Continue to populate new GFS .25 degree output and NCEP HRRR output to an S3 bucket for access. We did not place a TDS on this collection as this output is available from our standard sources.
- Unidata continues to get asked, from other UCAR/NCAR groups, to partner and lend assistance in cloud computing...especially in the AWS cloud.

Product Generation for IDD

For the past three years, Unidata generated products for the IDD, FNEXRAD and UNIWISC data streams have been created by a VM hosted in the Amazon cloud. This production generation has been proceeding very smoothly with almost no intervention from Unidata staff.

CloudIDV, CloudStream, Cloud Control

- We have released the technology enabling CloudIDV in a form that can be easily leveraged by other projects looking to bring legacy software to the Cloud. We are currently trying to build the CloudStream community via conference presentations and outreach.
- In addition, we continue to experiment with CloudIDV on the Jetstream Cloud. We are investigating CloudIDV for data-proximate visualization of the WRF-hydro modeling system.

Open Commons Consortium Award

The Open Science Data Cloud, a resource of the Open Commons Consortium (OCC), provides the scientific community with resources for storing, sharing, and analyzing terabyte and petabyte-scale scientific datasets. The OSDC is a data science ecosystem in which researchers can house and share their own scientific data, access complimentary public datasets, build and share customized virtual machines with whatever tools necessary to analyze their data, and perform the analysis to answer their research questions. Unidata is a beta user of resources in the Open Science Data Cloud ecosystem and we have been provided cloud-computing resources on the Griffin cloud platform. Our allocations are renewed on a quarterly basis and Unidata is partnering with OCC on the NOAA Big Data Project. Given the limited staff resources and many ongoing cloud activities on AWS, Azure, and XSEDE environments, Unidata's activities on the OSDC have been in a temporary hiatus. We are hoping to ramp up our OSDC efforts in the upcoming months.

New Activities

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

CloudIDV, CloudStream, Cloud Control

Forthcoming Presentations

"Data-Proximate Analysis and Visualization in the Cloud using Cloudstream, an Open-Source

Application Streaming Technology Stack”, 2017 AGU Fall Meeting | December 11–15, 2017 – New Orleans, LA USA

Unidata Science Gateway

We aim to collaborate with Jeremy Fischer at IU, XSEDE and Rich Signell to experiment with the "Zero to JupyterHub" project. The goal is to take advantage of cloud scalability for on-demand use in a classroom setting, for example, with technologies such as OpenStack and Kubernetes.

Forthcoming Presentations

- "Unidata Science Gateway on the XSEDE Jetstream Cloud", Gateways 2017 | October 23–25, 2017 – Ann Arbor, MI USA
- "Reducing Time to Science: Unidata and JupyterHub Technology Using the Jetstream Cloud", 2017 AGU Fall Meeting | December 11–15, 2017 – New Orleans, LA USA
- "A Unidata JupyterHub Server: An Online PyAOS Resource for Students and Educators", 2018 AMS Annual Meeting | January 7–11, 2018 – Austin, TX USA

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

Unidata Science Gateway

We would like to promote and advertise the science gateway (<http://science-gateway.unidata.ucar.edu/>) to our community.

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

Unidata Transitioning to the Cloud

In the long-term, we would like to explore the possibility of migrating some core Unidata services onto the cloud.

Areas for Committee Feedback

We are requesting your feedback on the following topics:

1. What clouds are our community using, either commercial (e.g., Amazon) or non-for-profit (e.g., NSF XSEDE Jetstream)?
2. What new cloud technologies are our community using and investigating on their own initiative?
3. Who would like to volunteer to beta test CloudIDV?
4. Who would like to volunteer to beta test: <https://jupyter-jetstream.unidata.ucar.edu>

Relevant Metrics

Docker image downloads are available from [Unidata's Dockerhub repository](#).

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

- 1. Enable widespread, efficient access to geoscience data**
Making Unidata data streams available via various commercial (e.g., Amazon) and not-for-profit (e.g., NSF XSEDE) cloud services will allow our community to access data quickly and at low or even no cost. Moreover, our users can benefit from high data bandwidth capability provided by various cloud computing platforms, and in some cases, Internet2 capability. Lastly, cloud computing offers the possibility of accessing geoscience data in a "data-proximate" manner where users can perform analysis and visualization on, at times, unwieldy data sets next to where the data reside.
- 2. Develop and provide open-source tools for effective use of geoscience data**
Containerization technology complements and enhances Unidata technology offerings in an open source manner. Unidata experts install, configure and in some cases, security harden Unidata software in containers defined by Dockerfiles. In turn, these containers can be easily deployed on cloud computing VMs by Unidata staff or community members that may have access to cloud-computing resources. Unidata staff develop Docker containers in an open-source manner by employing software carpentry best-practices and distributed version control technology such as git.
- 3. Provide cyberinfrastructure leadership in data discovery, access, and use**
Unidata is uniquely positioned in our community to experiment with cloud computing technology in the areas of data discovery, access, and use. Our efforts to determine the most efficient ways to make use of cloud resources will allow community members to forego at least some of the early, exploratory steps toward full use of cloud environments.
- 4. Build, support, and advocate for the diverse geoscience community**
Transitioning Unidata technology to a cloud computing environment will increase data availability to new audiences thereby creating new and diverse geoscience communities.

Prepared October 2017

Status Report: Community Services

April 2017 - October 2017

Doug Dirks, Jeff Weber, Joshua Young

Activities Since the Last Status Report

News@Unidata blog

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

- [Python-Focused Software Training Workshop at the University of Wisconsin](#)
- [Welcome Communications Intern Aubrey Kane](#)
- [Welcome Summer Intern Tyler Wixtrom](#)
- [Get Involved in Python Development for the Geosciences!](#)
- [2017 Community Equipment Awards](#)
- [Cloud-Based Visualization at Pennsylvania State University](#)
- [Welcome SOARS Intern Briah' Davis](#)
- [MetPy Mondays!](#)
- [Registration Open for MetPy Short Course at 2018 AMS Annual Meeting](#)
- [Solar Eclipse Visualizations](#)
- [New: Contributor License Agreement for Unidata Open Source Projects](#)
- Software release information
- Community job postings
- Community meetings and other announcements

Dependencies, challenges, problems, and risks include:

- Finding community members willing to contribute stories (or story ideas) for the blog is an ongoing challenge

Community Outreach and Services

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

Progress has been made on the following:

- Fine tuning the text of the next Unidata Strategic Plan, with guidance from the Strategic Advisory Committee.
- Work with Communications Intern Aubrey Kane on evaluating additional freely-available open source tools for inclusion in the DMRC.
- Planning for participation in upcoming professional society conferences (2017 AGU Fall Meeting and 2018 AMS Annual Meeting)
- Planning and support for the 2018 Users Committee Triennial Workshop
- Engagement with CUAHSI to support the NFIE and WRF-Hydro at the NWC

- Engagement with EarthCube as a Co-Chair of the Liaison Team
- Continue to serve on the CUAHSI HIS standing committee
- We continue to update Unidata's social media channels (Facebook, Twitter, Google+)
- We continue to publish short videos/screencasts on the [Unidata YouTube channel](#).
- We continue to actively support the NCAR/SOARS program
- Actively participate in Super Science Saturday
- Engage and support the Undergraduate Leadership Workshop (ULW) at UCAR
- Support the development and operation of the UCAR:NCAR Equity and InclusiON (UNEION) community of practice

Dependencies, challenges, problems, and risks include:

- Facilitating community adoption of new technological services (cloud, etc)
- Engagement with Unidata social media streams among community members is not particularly high.

Ongoing Activities

We plan to continue the following activities:

- Engagement with EarthCube, ESIP, and science or cyber communities at large
- NAWIPS migration to AWIPS, including the overall AWIPS project
- Ongoing development of news articles for publication through News@Unidata
- Continue to support and contribute to governing committees
- Seminars
- Outreach
- Inclusion and equity
- Engagement with professional societies
- Support for cloud-related projects
- Further development of the Data Management Resource Center
- Further work on Agile Data Curation
- Site visits as the budget allows
- Engage other UCAR/NCAR divisions regarding Unidata software use i.e. CESM/IDV

New Activities

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

- Expanded emphasis on cloud-related activities
- Seek additional resources to continue the online training resource focused on Python and Unidata services and tools
- Support the Users Committee's work towards planning the next Triennial

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

- Continue to engage the hydrologic community regarding WRF-Hydro/IDV interactions

- and the National Water Center's efforts
- Seek additional opportunities to engage and listen to the community

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

- Provide additional data management and cloud-related training

Areas for Committee Feedback

We are requesting your feedback on the following topics:

1. Please visit the [DMRC](#) and [OTP](#); these are both pilots but we would welcome your suggestions.

Relevant Metrics

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

All community pages

Most recent six months:

- 56,470 unique pageviews (51,121 in previous period)
- 10.1% of total unique pageviews (8.4% in previous period)

Top community pages

1. All blog pages
41969 unique pageviews (39592 in previous period)
74% of total community pageviews (77% in previous period)
2. www.unidata.ucar.edu/events
8476 unique pageviews (4045 in previous period)
15% of total community pageviews (8% in previous period)
3. www.unidata.ucar.edu/community
2782 unique pageviews (3814 in previous period)
5% of total community pageviews (7% in previous period)
4. www.unidata.ucar.edu/about
2660 unique pageviews (2971 in previous period)
5% of total community pageviews (6% in previous period)

Social media statistics, March 31, 2017

1. # of Twitter followers: 738 (up from 631)

2. # of Facebook followers: 707 (up from 645)

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
We monitor and collaborate with data sources to stay apprised of impending changes and to advocate for the needs of our user community. We provide user workshops, tutorials, and community workshops to help build supportive relationships between community members.
2. **Develop and provide open-source tools for effective use of geoscience data**
We promote Unidata tools and software for multi-disciplinary use, with an eye toward finding additional research and educational communities that can benefit from our work.
3. **Provide cyberinfrastructure leadership in data discovery, access, and use**
We work with government and industry data providers to secure access to data for Unidata community members.
4. **Build, support, and advocate for the diverse geoscience community**
We coordinate with our governing committees to find ways to expand Unidata's community participation. We use our web site, electronic newsletters, and social media to keep community members informed about enhanced data services, software tools, and cyberinfrastructure.

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

Prepared September 2017

Status Report: Community Equipment Awards

Sponsored by the National Science Foundation

April 2017 - October 2017

Admin Group

The NSF provides the Unidata Program Center up to \$100k in equipment grant funds each year. In alignment with the Unidata 2018 proposal, the Equipment Awards Program is designed to broaden participation and promote the use of Unidata tools and systems (e.g., THREDDS, NetCDF, IDV, GIS connections) to support education and research on various aspects of climate studies (e.g., diagnostics, change and impacts), by providing grants to be used in the procurement of new computers and equipment including upgrades to existing classroom and laboratory equipment.

This year, special consideration was given to proposals that included one or more of the following:

- Production of online training materials (code notebooks, video tutorials, online documentation, or similar resources) that can be shared with the Unidata community
- Projects that advance data-proximate analysis of large remote datasets (e.g. GOES-R satellite data)
- Installation of equipment that provides student access to modern visualization tools such as AWIPS, IDV, or Jupyter notebooks

A Request for Proposals was sent out on January 20, 2017 with a March 31, 2017 submission deadline. The Review Panel met on April 20 at the Unidata Program Center and recommended that the following proposals be funded:

- Central Michigan University, John Allen, "Enhancing Undergraduate Python and Modeling Skills: A Jupyter Notebook Multi-Core Server at Central Michigan University" [proposal](#)
- College of DuPage, Victor Gensini, "Eyes in the Sky: Installation of a GOES-16 Data Server at College of DuPage" [proposal](#)
- Embry-Riddle Aeronautical University, Chris Herbster, "Infrastructure Enhancements to Support Educational Use of GOES-16" [proposal](#)
- University of Hawaii, Steven Businger, "Pacific THREDDS Data Server" [proposal](#)
- University of Miami, Brian Mapes, "University of Miami Data Server 2017" [proposal](#)

Congratulations to all of the recipients and a special thank you to the Review Panel and the NSF for making the Unidata Equipment Awards program possible.

Areas for Committee Feedback

We are requesting your feedback on the following topics:

1. Possible theme(s) for the 2018 Unidata Community Equipment Awards;
2. Please consider volunteering to serve on the 2018 Review Panel;
3. Suggestions from previous panel members on how to improve the program

Relevant Metrics

Since taking over the management and administration of the Equipment Awards program in 2003 on behalf of the NSF, Unidata has made 92 awards totaling close to \$1,200,000.

Prepared September 2017

Status Report: GOES-R

April 2017 - October 2017

Mike Schmidt, Jeff Weber, Tom Yoksas

Activities Since the Last Status Report

- Test IDD distribution of GOES-16 imagery and space weather products. NCAR/RAL was the first site to use this service to provide GOES-16 data for their science objectives
- Install robust data serving platforms to be used to provide programmatic access to GOES-16 ABI imagery, space weather products and the GLM at the Mesa Lab (new)
- Continue investigating implementing UW/SSEC's "fanout server" (redistribution of the GRB-200 UDP unicast stream over TCP) via an Linux Virtual Server (LVS) cluster
 - This would provide high availability for the GRB-200 output which could then provide real-time GRB input for sites other than UCAR
- Secure funding to install GOES-S ingest and data serving capabilities. Monies for the downlink portion of the installation have already been received. We need to pursue funding for the data processing and serving machines.

Ongoing Activities

We plan to continue the following activities:

- Ingest the GOES ReBroadcast (GRB) from GOES-16 in real-time using the 4.5 m satellite dish that we installed on the eastern pad at the NCAR Mesa Lab using the NOAA funded UW/SSEC/CIMSS Community Satellite Processing Package (CSPP) for Geostationary Data (GEO) package
 - CSPP GEO is responsible for GRB ingest and productization (e.g., stitching together of broadcast tiles)
- Upload all ABI imagery, Space Weather products and Geostationary Lightning Mapper (GLM) to an S3 bucket in Amazon Web Services (AWS)
 - Transfer of ABI imagery and space weather products has been active since the first day of GRB broadcast using storage resources provided free-of-charge by Amazon
- Continue working with NCAR/EOL to support their CSPP GEO installation at the NCAR Mesa Lab
 - This installation will provide an in-house redundancy for GOES-16 ingest
- Reach out the greater Unidata community to learn about their GOES-16 data needs
- Investigate additions to the IDD CONDUIT data stream that would be useful for creation of new GOES-16 based Level 2 products

Future Activities

- Serve GOES-16 Advanced Baseline Imager (ABI) imagery via McIDAS ADDE and the TDS
 - Testing is already in progress, but has not been announced to the community since NOAA has not yet declared and end to testing/modifications of the GOES-16 imagery and products being sent in the GRB and NOAAport

- Announce an LDM/IDD feed of ABI imagery, space weather products and the GLM
 - We previously commented that the expected data volume of GOES-16 data would be too large to relay in the IDD. Our opinion was changed when the data volume was drastically reduced by the turning on of netCDF4 compression of the data being received in the GRB
- Investigate approaches that would insulate GRB ingestion from long standing twice per year power downs in the NCAR Mesa Lab facility
- Unidata-Wisconsin (UNIWISC) IDD imagery will be updated to incorporate GOES-16 imagery. It is likely that the spatial resolution of several of the products will be increased during this update (ETA end of 2017, beginning of 2018)

Relevant Metrics

- Lots O Data!

The volume of GOES-16 imagery, space weather and GLM products can be seen in the real-time statistics plot from our ingest machine:

http://rtstats.unidata.ucar.edu/cgi-bin/rtstats/iddstats_vol_nc?EXP+ingest.unidata.ucar.edu

The volume of ABI imagery and GLM Level 2 products can be seen in the real-time statistics plot from NCAR/RAL's ingest machine:

http://rtstats.unidata.ucar.edu/cgi-bin/rtstats/iddstats_vol_nc?EXP+satops3.rap.ucar.edu

Questions for Committee Members

- What new coverages and increased spatial and temporal resolutions would should be considered for UNIWISC IDD feed imagery?
- What kind(s) of data access methods are most desired/usable for the community?
- Other?

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
Standing up ADDE and TDS data services for real-time GOES-R/S data will benefit the greater Unidata community.
2. **Develop and provide open-source tools for effective use of geoscience data**
The IDD is powered by the Unidata LDM-6 which is made freely available to all. The Unidata NOAAPort ingest package is being used by a variety of university and non-university community members. Both the LDM and NOAAPort ingest packages are bundled in AWIPS.
3. **Provide cyberinfrastructure leadership in data discovery, access, and use**

The community-driven IDDs provide push data services to users an ever increasing community of global educators and researchers.

- 4. Build, support, and advocate for the diverse geoscience community**
Providing access to data in real-time is a fundamental Unidata activity.

Prepared *October 2, 2017*

Status Report: Internet Data Distribution

April 2017 - October 2017

Mike Schmidt, Jeff Weber, Steve Emmerson, Tom Yoksas

Activities Since the Last Status Report

Internet Data Distribution (IDD)

IDD data volumes continue to increase. The following output is from a Linux-based motherlode clone that the UPC operates on behalf of the community, lead.unidata.ucar.edu:

Data Volume Summary for lead.unidata.ucar.edu

Maximum hourly volume 72480.665 M bytes/hour

Average hourly volume 49696.601 M bytes/hour

Average products per hour 390913 prods/hour

Feed	Average (M byte/hour)		Maximum (M byte/hour)	Products number/hour
FSL2	14651.912	[29.483%]	17446.562	22399.975
CONDUIT	7887.856	[15.872%]	21436.385	87792.900
NGRID	7596.554	[15.286%]	13110.295	40235.375
EXP	6085.977	[12.246%]	8779.037	3440.750
NEXRAD2	5387.548	[10.841%]	7704.114	62359.450
NOTHER	3699.461	[7.444%]	7465.859	7456.375
NEXRAD3	1803.479	[3.629%]	2401.904	90171.375
FNMOC	1214.143	[2.443%]	4750.510	3301.400
HDS	749.161	[1.507%]	1156.088	27504.625
GEM	203.950	[0.410%]	1157.301	1229.975
NIMAGE	156.532	[0.315%]	349.965	180.950
FNEXRAD	123.776	[0.249%]	153.917	103.850
IDS DDPLUS	68.174	[0.137%]	81.383	44113.425
UNIWISC	64.426	[0.130%]	120.964	46.875
LIGHTNING	3.654	[0.007%]	9.726	575.700

Ongoing Activities

We plan to continue the following activities:

- Unidata continues to receive the NOAA/GSD experimental High Resolution Rapid Refresh (**HRRR**) grids (both 2D and 3D fields) in an LDM/IDD feed from NOAA/GSD and feed these products to a small number of university sites on **hrrr.unidata.ucar.edu**.
- The NCEP operational HRRR is being served with other model output at:

<http://thredds-jumbo.unidata.ucar.edu/thredds/catalog/idd/forecastModels.html>____
(.xml for machine access)
- Other data sets we continue to explore with NOAA/GSD/ESRL are:
 - [FIM](#)
 - [HIWPP](#)
- NCEP (operational) HRRR fields and forecasts times have been added to the CONDUIT IDD datastream.

NOAAPort Data Ingest

- Ingest of the upgraded NOAAPort Satellite Broadcast Network (SBN) products and their relay to end-users via the IDD has been “operational” at the UPC since the August 2014.
- The NOAAPort-derived data streams (HDS, IDS|DDPLUS, NGRID, NIMAGE, NEXRAD3 and NOTHER) are being redundantly injected into the IDD at three geographically separate locations: Unidata, UW/SSEC, and LSU/SRCC. The NOTHER data stream contains GOES-16 tiles that need to be stitched together to make full image scenes usable to end-user applications.

We continue to look for a fourth ingest site to increase robustness of the IDD distribution of NOAAPort derived data.

- Unidata's NOAAPort ingest package is bundled with current versions of the LDM. The current LDM release is v6.13.6.
- Raytheon continues to submit modifications to the LDM for enhanced AWIPS functionality.

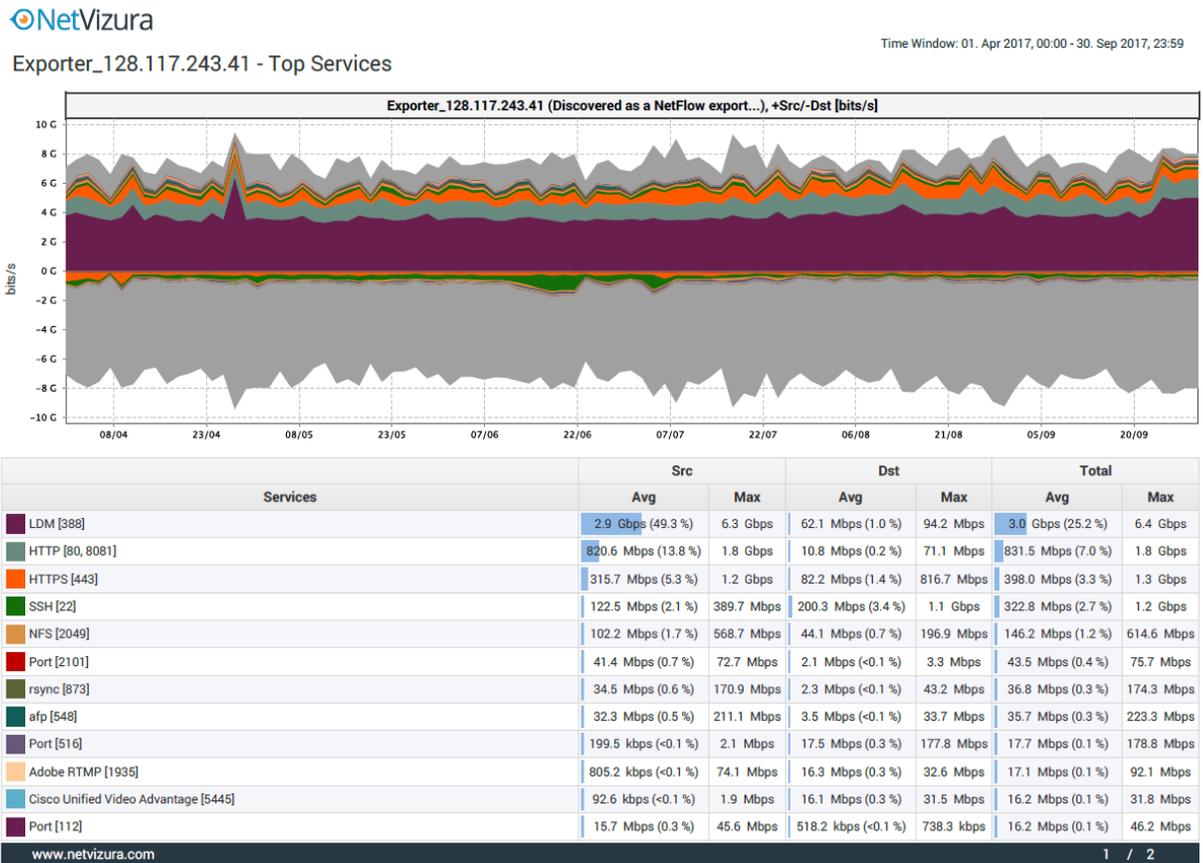
Relevant Metrics

- Approximately **590** machines at **250** sites are running LDM-6 **and** reporting real time statistics to the UPC. Unidata staff routinely assist in the installation and tuning of LDM-6 at user sites as a community service.

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

The IDD relay cluster, described in the June 2005 CommunitE-letter article Unidata's IDD Cluster, routinely relays data to more than 1250 downstream connections.

Over the period from April 1 through September 30, 2017 the average volume of LDM/IDD data flowing from the UCAR/NCAR network averaged around 3 Gbps (~31 TB/day), and peak rates reached 6 Gbps (which would be ~63 TB/day if the rate was sustained).



Cluster real server backends and accumulator nodes routinely have instantaneous output volumes that can exceed 1.5 Gbps. Bonding of pairs of Ethernet interfaces was needed to be able to support these output data rates. The next generation of cluster machines will need to have 10 Gbps Ethernet capability.

The 6-8 GB/hr increase in IDD data volume shown in the volume list for **lead.unidata.ucar.edu** reflects the test relay of GOES-16 data.

Questions for Committee Members

- Suggestions regarding content of data streams like CONDUIT, FNEXRAD, UNIWISC and NLDN LIGHTNING? ... We (UPC, the Unidata community and UAlbany for the NLDN component of LIGHTNING) have control of the content of these data streams, so they are open for suggestions.
- We have been preparing the top level IDD relay clusters we maintain to support relay of the full volume of GOES-16 imagery and products being received on the GOES-R downlink that we installed at the NCAR Mesa Lab. The data volume for this future feed can be seen in the EXP feed in the volume listing for lead.unidata.ucar.edu. The question for the governing committees is if relay of the GOES-16 data would be desired by the community, **and** if end users can handle the data volume.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
A project like the IDD demonstrates how sites can employ the LDM to move data in their own environments.
2. **Develop and provide open-source tools for effective use of geoscience data**
The IDD is powered by the Unidata LDM-6 which is made freely available to all. The Unidata NOAAPort ingest package is being used by a variety of university and non-university community members. Both the LDM and NOAAPort ingest packages are bundled in AWIPS.
3. **Provide cyberinfrastructure leadership in data discovery, access, and use**
The community-driven IDDs provide push data services to users an ever increasing community of global educators and researchers.
4. **Build, support, and advocate for the diverse geoscience community**
Providing access to data in real-time is a fundamental Unidata activity. The IDD-Brasil, the South American peer of the North American IDD operated by the UPC, is helping to extend real-time data delivery outside of the U.S. to countries in South America and Africa. The Universidad de Costa Rica is actively pursuing IDD relay in the Latin America based IDD-Caribe.

Prepared October 2, 2017

Status Report: IDV with RAMADDA

April 2017 - October 2017

Yuan Ho, Julien Chastang

Activities Since the Last Status Report

IDV System Changes

__Latest netCDF-Java Version__

The version of the netCDF-Java library currently distributed with the latest stable IDV version (5.4) is the 4.6.10 (April 19, 2017). The prior version of netCDF-Java to be distributed with the IDV was 4.6.7. There have been many improvements and bug fixes in that range. [The complete release notes for these versions can be found here.](#)

__ISL Changes__

- Improved ISL offscreen mode
- Fixed ISL time driver start and end time bug
- Various improvements related to saving and loading of bundles with respect to time and region matching
- New ISL feature to publish to the RAMADDA

__Java Version__

The IDV is distributed with Java 8 u51. Because of a Windows 10 related bug on some older Intel hardware, this will be the latest version of Java we will be deploying with the IDV for some time.

__Java 3D Version__

The IDV is distributed with Java 3D version 1.6.0 Final. Java 3D 1.7 prerelease is under development and continues to be supported in some capacity by the open source community.

__HTTP to HTTPS__

As US government servers transition from HTTP to HTTPS, we update URL references in the IDV (rbi files, plugin files, etc.) to reflect these changes.

__Derived Quantities__

- Relative humidity calculation improvements for handling supersaturation.
- New Jython grid diagnostic formulas for calculating the wind shear between discrete layers.

__Point Observations__

Can now display relative humidity, wind chill, and heat index for metar point observations

__Plugins__

The IDV Jetstream plugin allows easy access to Jetstream installations of the TDS, RAMADDA and ADDE from the IDV. The plugin URL is `http://science-gateway.unidata.ucar.edu/jetstream.jar`

IDV Display Changes

__GOES-16 GLM (Global Lightning Mapper) data in the IDV__

A new GLM plugin has been developed to add three new data source types in the IDV (5.4 and nightly): GLM GROUPS Data Files, GLM EVENTS Data Files, and GLM FLASHES Data Files. These three new data source types will enable IDV users to load and display GLM data from Local disk and remote TDS systems.

__Super High Resolution Display__

Added a new FAQ item to help IDV users to deal with the small font and display issues in a super high Resolution screen.

__2D/3D Trajectories__

Improvements and enhancements in the underlying VisAD library has enabled a new grid 2D/3D trajectory feature in the IDV. The 2D/3D grid trajectories can be colored by speed or other parameters. The initial area of the trajectory is the entire domain of the dataset. However, regions can be selected with points, rectangles and closed polygons.

There are a variety of grid trajectory types including ribbon, cylinder, deform ribbon, point.

__Latest Version of VisAD__

The SSEC team at UW, Madison has made a number of improvements to support 3D trajectories.

__Skew T__

Fixed a bug concerning CAPE and CIN calculations in the Skew-T display.

__Transect Line__

Fixed the problem of the transect line modification in the -180 to 0 longitude domain.

__IDV Web Cams__

The IDV webcam list continues to be updated regularly

IDV WRF-Hydro Collaboration

Jeff Weber and Yuan are collaborating with David Gochis, and Dave Yates from NCAR-RAL assisting them ensuring WRF-Hydro data is CF compliant regarding Point Feature type. In addition, we are helping them with visualization of this dataset in the IDV.

IDV Release

IDV [5.4](#) was released in May of 2017.

IDV EarthCube Proposal Awarded

In collaboration, with University of Miami Professor Brian Mapes, Unidata submitted an EarthCube proposal: "Drilling down from a statistics scatterplot to pre-populated case notebooks". This proposal was awarded \$1.3 million over three years.

Unidata hires Celia Hedrick, a Software Engineer II to work on this project.

IDV Publication Highlights

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)).

RAMADDA and ADDE in the Cloud

Building upon previously described Docker containerization efforts, [we have deployed RAMADDA on the XSEDE Jetstream cloud](#). This installation of RAMADDA has a ["view" of the IDD on Jetstream](#). In addition, there is [a collection of IDV bundles](#) that retrieve data from the TDS installation on Jetstream.

ADDE is a data server developed by SSEC that primarily serves satellite and radar imagery and point observations to the IDV. An ADDE server is available at `adde-jetstream.unidata.ucar.edu` over (the usual) port 112.

IDV and RAMADDA Training, Conference Attendance and Presence

__Gateways 2017__

- Unidata will attend the Gateways 2017 conference in Ann Arbor. [Our efforts on deploying RAMADDA, ADDE and its use for the IDV will be described in a conference poster.](#)

__2017 American Geophysical Union (AGU) Annual Meeting__

- Yuan will attend 2017 AGU and present: The Comparison of Point Data Models for the Output of WRF Hydro Model in the IDV

__2018 American Meteorological Conference (AMS) Annual Meeting__

- Yuan will attend 2018 AMS and present: The new trajectory displays in the UNIDATA's IDV

Ongoing Activities

We plan to continue the following activities:

__IDV Instructional Videos__

No instructional videos on the IDV are planned at this time. We would appreciate input and suggestions on specific video topics.

__IDV Seam Issues__

We will continue to collaborate with the Unidata netCDF-Java team to devise a solution for long-standing longitude seam issues in the IDV.

__Investigation of Java 3D Alternative__

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

New Activities

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

__GOES-16 GLM Collaboration__

At the 2017 EarthCube netCDF-CF Workshop, Unidata began discussions with SSEC, Randy Horne (Excalibur labs), Professor Eric Bruning (Texas Tech) on a path forward for displaying GOES-16 GLM (Global Lightning Mapper) data in the IDV.

Areas for Committee Feedback

We have no questions at this time.

Relevant Metrics

__E-Support__

The IDV team continues to provide the geoscience community with high-quality support through e-support software and idv-users maillist. The volume of e-support remains high and constitutes a large fraction of our daily activities. In the last half year the IDV team has closed ~90 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.

__Usage Metrics__

Raw IDV usage metrics, are available here

<http://www.unidata.ucar.edu/software/idv/logging/left.html>.

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 111 “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 13,200 times compared with 11,500 from six months ago.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**

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

2. **Develop and provide open-source tools for effective use of geoscience data**

The IDV has been an open-source project for several years. The IDV is available on the github version control platform for greater open-source collaboration. The IDV provides

users the unparalleled ability to analyze, integrate, and visualize heterogeneous geoscience data in two, three, and four dimensions. The IDV coupled with RAMADDA enables geoscience specialists the capability to share and collaborate their IDV analysis via social scientific networks.

3. **Provide cyberinfrastructure leadership in data discovery, access, and use**

RAMADDA allows geoscience specialists the ability to search and publish their IDV bundles on-line. Unidata's RAMADDA installation enables the IDV team to communicate more effectively to our users concerning their IDV issues. Specifically, during support ticket conversations, the IDV team requests that users upload pertinent data to RAMADDA for analysis. One of RAMADDA's best features is the ability to upload a CDM file and obtain the OpenDAP link from the new entry. The DAP link can be shared and opened in the IDV. RAMADDA also has "server-side view" capability where a specific part of the file system can be made available through the RAMADDA interface. This feature can be helpful to view LDM data feeds, for example. The IDV team also takes advantage of RAMADDA to share instructional IDV screencasts with users.

4. **Build, support, and advocate for the diverse geoscience community**

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

Status Report: International Activities

April 2017 - October 2017

Tom Yoksas, Yuan Ho

Activities Since the Last Status Report

Renewed effort aimed at creating an IDD-Caribe

Dr. Marcial Garbanzo Salas of the Universidad de Costa Rica is leading an effort aimed at making an IDD-Caribe a reality. Unidata is playing a supporting role in this effort.

New/Planned Activities

No new activities are currently planned

Ongoing Activities

- Imagery from UCAR GOES East/West ingest systems continue to be routinely accessed by international users in North, Central and South America using McIDAS-X, IDV, and McIDAS-V.
- Even though access to GOES-16 imagery and products has not been announced (we are waiting for GOES-16 data to be declared operational), we are seeing substantial access to the imagery via ADDE on our ingest machine.
- Use of Unidata tools, especially netCDF, the IDV and AWIPS continues to grow internationally. Use of GEMPAK internationally seems to have plateaued.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

- 1. Develop and provide open-source tools for effective use of geoscience data**
The majority of tools downloadable from Unidata are available free-of-charge to everyone (the exception being McIDAS-X).
- 2. Provide cyberinfrastructure leadership in data discovery, access, and use**
Activities of the Unidata Program Center are routinely provided to the worldwide atmospheric science community. Strategic partnerships with leading organizations in other countries minimize the impact on UPC staff.
- 3. Build, support, and advocate for the diverse geoscience community**
By informing the international atmospheric science community of the products, data and services available in the Unidata Program, an extended community has been enabled. Non-U.S. users of products available from Unidata reflect, in a number of cases, minority constituencies in the U.S. atmospheric science community.

Prepared *October 2, 2017*

Status Report: LDM

April 2017 - October 2017

Steve Emmerson, Tom Yoksas, Mike Schmidt

Activities Since the Last Status Report

LDM

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

Progress has been made on the following:

- Changes to accommodate Docker
- Changes to accommodate non-standard installations
- Improvements to the time-keeping subsystem
- Improvements to pqing(1):
 - Added XON/XOFF software flow control
 - Improved response to "ldmadmin stop"
 - Ensure conforming TTY behavior
- Improvements to pqact(1) documentation

Dependencies, challenges, problems, and risks include:

- Dealing with missing GEMPAK table entries represents a problem area
- The LDM is sometimes held responsible for decisions made by the NWS on how to categorize data products.

Multicast LDM (aka LDM-7)

The multicast LDM project is separately funded by CISE in NSF. The goal is to reduce the outgoing bandwidth requirement of the LDM -- yet retain the current level of reliability -- by converting it into a hybrid system that combines use of the new, semi-reliable multicast protocol developed at the University of Virginia with the time-tested unicast capability of the current LDM.

Another multi-year grant has been awarded by the NSF for this project. The funds were finally available at the end of September.

Ongoing Activities

We plan to continue the following activities:

- Support and maintenance of the LDM

- Continue adapting the LDM to the Docker container technology to support cloud activities and (perhaps) make life easier for LDM users

New Activities

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

- Stand-up an LDM-7 server at the Unidata Program Center
- Have the University of Virginia receive IDD traffic via LDM-7
- Stand-up LDM-7 clients at identified universities

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

- Continued development and deployment of LDM-7

Relevant Metrics

- Data on the LDM package can be found [here](#)
- The LDM system at the Unidata Program Center powers the IDD. Metrics on that program can be found in the IDD status report.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
By enabling researchers, teachers, and students to obtain a wide variety of meteorological and related data in near real time and at no cost via the Internet.
2. **Provide cyberinfrastructure leadership in data discovery, access, and use**
By using the LDM to move data into the cloud and developing multicast technologies.

Prepared *March, 2017*

Status Report: McIDAS

April 2017 - October 2017

Tom Yoksas

Activities Since the Last Status Report

Aside from routine updates/bugfixes to existing code and tables, the main area of investigation has to been to create ADDE servers for NOAAPort delivered GOES-16 satellite imagery.

Current Activities

- Unidata McIDAS version 2017 is being prepared for release. v2017 includes all SSEC versions up to and including the current release, v2017.1, Unidata development, updates and bugfixes, and will contain SSEC McIDAS v2017.2 code modifications and additions as soon as they are made available.

Changes to Unidata McIDAS continue to be made through an **addendum** process.

The latest releases feature the following:

- Updated ADDE servers for Himawari and GOES-R imagery
- ADDE server for GOES-R imagery: The core McIDAS ADDE servers do not support the tiled GOES-16 image sectors that are being distributed in the NOAAPort SBN. A new ADDE server has been developed to support the NOAAPort-delivered images. Development of this server was put on hold during September while the NWS tests sending of image sectors that have not been remapped into conical projections. It is anticipated that the server development will proceed quickly now that the NWS test is over.

Ongoing Activities

We plan to continue the following activities:

- SSEC McIDAS Advisory Committee (MAC)

The UPC (Yoksas, Ho) continues to participate as the Unidata representative to the McIDAS Advisory Committee (MAC) that is operated by SSEC.

The MAC was assembled by UW/SSEC to advise SSEC on McIDAS-X users needs/concerns/desires for development in the next generation McIDAS, McIDAS-V. The MAC was modeled after the Unidata IDV Steering Committee.

- Interest in McIDAS by Non-core Users

The UPC continues to receive requests for McIDAS-X and help using McIDAS-X from international university users, U.S. government agencies and other non-traditional

Unidata users (e.g., private businesses, etc.). Government agencies and non-traditional Unidata users are referred to UW/SSEC for access to McIDAS; international educational community user requests are granted on a case-by-case basis after they provide a clear statement of their acceptance of the terms of use provided by SSEC.

New Activities

Ongoing Activities

Continued support of existing and new community members.

New Activities

Add support for new types of data when they become available, otherwise McIDAS-X support is in maintenance mode.

Relevant Metrics

- Internet2 bandwidth usage by the Unidata McIDAS ADDE servers routinely exceeds 8 TB/month. The amount of data served by Unidata ADDE instances has dropped recently, but it is expected to increase substantially when GOES-16 data is made available.
- [McIDAS-X/-XCD Inquiry Metrics](#)

ldm-mcidas Decoders Activities

Development

ldm-mcidas releases are made when needed to support changes in software development and operating system environments. **ldm-mcidas** v2012 was released at the end of September, 2012. This package will be updated to support changes to various data streams in the December, 2017 time frame.

Geostationary Satellite Data Ingest and Data Serving

Unidata continues to ingest GOES-East and GOES-West imager data at the UCAR Foothills Lab campus in Boulder.

- Direct, programmatic access to real-time GOES-East (GOES-13) data via McIDAS ADDE routinely exceeds 3.2 TB/month.
- Direct, programmatic access to real-time GOES-West (GOES-15) data via McIDAS ADDE routinely exceeds 2.6 TB/month.
- Direct, programmatic access to real-time GOES-16 data via McIDAS ADDE has been exceeding 1.3 TB/month ***even though this service has not yet been announced.***

Planned Activities

Ongoing Activities

Continued ingest and serving of GOES-East and GOES-West imagery from the current constellation of GOES GVAR platforms. This effort requires sporadic maintenance of the satellite ingest and data serving equipment.

New Activities

Finish installing the data serving hardware for our new GOES-R downlink/data distribution facility at the NCAR Mesa Lab. This activity is imminent.

Continue moving GOES-R imagery and products to an S3 bucket in AWS in near real-time.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

- 1. Enable widespread, efficient access to geoscience data**
*MclIDAS remains **the** application of choice for the satellite meteorology community. The Advanced Data Distribution Environment (ADDE) component of MclIDAS was the first application offered by Unidata to provide remote, programmatic access to a wide variety of data that is important to the atmospheric science community.*
- 2. Develop and provide open-source tools for effective use of geoscience data**
The fifth generation of MclIDAS, MclIDAS-V, unlike its predecessors, is a fully open source application that is in wide scale and growing use in the worldwide satellite meteorology community. MclIDAS ADDE continues to evolve and provide access to increasing volumes of image and non-image data.
- 3. Provide cyberinfrastructure leadership in data discovery, access, and use**
Concepts articulated in ADDE inspired the development of THREDDS (to address the lack of rich metadata available in ADDE) and RAMADDA. ADDE remains one of the most used data services in the Unidata suite. ADDE servers in the SSEC Data Center are currently serving over 1 TB per day.
- 4. Build, support, and advocate for the diverse geoscience community**
MclIDAS is sought for use by those interested in satellite meteorology worldwide.

Status Report: netCDF

April 2017 - October 2017

Ward Fisher, Dennis Heimburger

Activities Since the Last Status Report

We are using GitHub tools for C, Fortran and C++ interfaces to provide transparent feature development, handle performance issues, fix bugs, deploy new releases and to collaborate with other developers. Additionally, we are using docker technology to run netCDF-C, Fortran and C++ regression and continuous integration tests. We currently have 67 open issues for netCDF-C, 15 open issues for netCDF-Fortran, and 13 open issues for netCDF-C++. The netCDF Java interface is maintained by the Unidata CDM/TDS group and we collaborate with external developers to maintain the netCDF Python interface.

In the netCDF group, progress has been made in the following areas since the last status report:

- Further extension of the netCDF build-and-test platforms using Docker technology.
- Further enhancements to the netCDF documentation.
- Extended continuous integration platforms have been adopted.
- An architecture roadmap is available describing how the netcdf-c library will support thread-safe operation in *nix* and Windows environments. The draft proposal is available [as netcdf-c github issue #382](#).
- Support for the DAP4 protocol is now part of the code-base. It has been verified for consistency against the Thredds Java DAP4 implementation. DAP4 remote testing is currently disabled until a new test server can be established. Our expectation is that the test server will be stood up on the Jetstream cloud.
- We have seen an uptick in the number of contributions to the netCDF code base(s) from our community. While these contributions require careful review and consideration, it is encouraging to see this model of development (enabled by our move to GitHub) being more fully embraced by our community.
- Increased support for native CDF5 on 64-bit platforms, potentially obviating the need for the parallel netcdf library for reading/writing CDF5 files in serial fashion.

Dependencies, challenges, problems and risks include:

- Small group (and shrinking) of developers for supporting large project.
- Dependency on HDF5, controlled by external group.
- Slow progress in user adoption of netCDF-4 features.

Ongoing Activities

We plan to continue the following activities:

- Provide support to a large worldwide community of netCDF developers and users.
- Continue development, maintenance, and testing of source code for multiple language libraries and generic netCDF utility programs.
- Improve organization of Doxygen-generated documentation for netCDF-C and Fortran libraries.
- Extend collaboration as opportunities arise, for increasing the efficiency of parallel netcdf-3 and netcdf-4.

New Activities

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

- Seek out, and prepare material for upcoming, conferences and other outreach opportunities.
- Work on reducing the defects reported by static analysis.
- Release the next versions of netCDF-C, netCDF-Fortran, netCDF-C++.
- Modernize the netCDF documentation to provide easy access to documentation for older versions of netCDF.
- Provide thread-safety for the netCDF C library.

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

- Continue integration of the upcoming ExaHDF5 features into the netCDF-C, Fortran and C++ interfaces.
- Release an official Windows port of the netCDF-Fortran and netCDF-C++ interfaces.
- Participate in development of new CF 2.0 conventions for climate and forecast simulation output and observational data in netCDF-4 form.
- Continue to encourage and support use of netCDF-4's enhanced data model by third-party developers.
- Create and release online educational material in the form of Youtube video tutorials for using netCDF.

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

- Implement support for Amazon S3 in the netCDF C library.
- Improve scalability to handle huge datasets and collections.
- Improve the efficiency of parallel netcdf3 and parallel netcdf4.

Areas for Committee Feedback

We are requesting your feedback on the following topics:

1. **To what extent is Amazon S3 used within your organization? Would you benefit from native netCDF support for S3 storage?**
2. **Are there any emergent avenues (stack overflow, etc) for user support which the netCDF team should investigate?**
3. **How can we encourage more user testing of the release candidates we provide?**

Relevant Metrics

There are currently about 183,700 lines of code (up from 142,810 lines of code) in the netCDF C library source. The Coverity estimate for defect density (the number of defects per thousand lines of code) in the netCDF C library source has been increased slightly from **0.36** six months ago to **0.7** today. According to Coverity static analysis of over 250 million lines of open source projects that use their analysis tools, the average defect density with 100,000 to 500,000 lines of code is **0.50**. The jump in defect density is a result of the addition of the **DAP4** code. As this is new code, the initial defects are still being worked out.

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

Currently, Google hits, for comparison, are:

- **889,000** for netCDF-3
- **930,000** for netCDF-4
- **464,000** for HDF5
- **89,700** for GRIB2

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

- **282** for netCDF-3
- **526** for netCDF-4
- **9,420** for HDF5
- **792** for GRIB2

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
by developing netCDF and related cyberinfrastructure solutions to facilitate local and remote access to scientific data.

- 2. Develop and provide open-source tools for effective use of geoscience data**
by supporting use of netCDF and related technologies for analyzing, integrating, and visualizing multidimensional geoscience data; enabling effective use of very large data sets; and accessing, managing, and sharing collections of heterogeneous data from diverse sources.
- 3. Provide cyberinfrastructure leadership in data discovery, access, and use**
by developing useful data models, frameworks, and protocols for geoscience data; advancing geoscience data and metadata standards and conventions; and providing information and guidance on emerging cyberinfrastructure trends and technologies.
- 4. Build, support, and advocate for the diverse geoscience community**
by providing expertise in implementing effective data management, conducting training workshops, responding to support questions, maintaining comprehensive documentation, maintaining example programs and files, and keeping online FAQs, best practices, and web site up to date; fostering interactions between community members; and advocating community perspectives at scientific meetings, conferences, and other venues.

Prepared *October 2017*

Status Report: Outreach Activities

April 2017 - October 2017

Ethan Davis, Sean Arms, Jeff Weber, & Josh Young

Activities Since the Last Status Report

[Open Geospatial Consortium \(OGC\)](#) and [Ocean Data Interoperability Platform \(ODIP\)](#) activities

Continue to work with OGC to augment international CF-netCDF standards that have been established over the last several years and on phase 2 of ODIP.

Progress has been made on the following:

- Represent Unidata at OGC Technical Committee meetings
- Chair [OGC NetCDF Standards](#) Working Group
- Serve as UCAR business and technical representative to the OGC.
- Start transfer of responsibility to other staff for ongoing projects following Ben Domenico's retirement

Dependencies, challenges, problems, and risks include:

- There is a question of what resources should be committed to these projects beyond the beginning of the 2017 calendar year when Ben retires from Unidata.

[EarthCube Cyberconnector Project](#)

Collaborative project with George Mason University to make Unidata real time datasets available to researchers and educators in other disciplines

Progress has been made on the following:

- Maintained periodic contact with Liping Di, the GMU PI on the project.
- Did mid-course correction to objectives based on one year no cost extension to the project.

Dependencies, challenges, problems, and risks include:

- The one year extension enables us to use the remaining resources available in the project to accomplish tasks that are in line with the strategic goals for the new version of the TDS and to ensure that our data are available via standard interfaces such as those provided by the Data Discovery and Access Broker of the Italian CNR groupPython Workshop Notebooks on alternative platforms.

Due to reduced FTE commitment, no effort is being expended on cloud software developmen

Ongoing Activities

- Coordination and collaboration with [NCAR GIS](#) (e.g. advising the GIS Program on cloud workflows)- Josh Young and Jeff Weber
- Represent Unidata and UCAR in OGC - Ethan Davis
- Represent Unidata in ESIP - Ethan Davis
- Represent Unidata at CUASHI - Jeff Weber
- Continue as co-PI on [Earthcube Cyberconnector project](#) - Sean Arms
- Participate in ODIP 2 as resources allow - Mohan Ramamurthy
- Participate in the GOES-16, Pathfinder, and VIIRS data working groups at AWS - Jeff Weber

Note that the responsibility for these ongoing activities is being transferred to other staff.

New Activities

Explore CloudIDV product generation at the National Water Center

Participate in [PREP](#) (Partnership for Resilience and Preparation)

Areas for Committee Feedback

None at this time.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
Work with representatives of other disciplines and serve on their governing boards where appropriate, e.g., NCAR GIS for Geographic Information Systems, CUAHSI (Consortium of Universities for Advancement of Hydrological Science), ODIP (Ocean Data Interoperability Platform), etc.
2. **Provide cyberinfrastructure leadership in data discovery, access, and use**
Continue to work with Open Geospatial Consortium (OGC) to augment international CF-netCDF standards that have been established over the last several years.
3. **Build, support, and advocate for the diverse geoscience community**
Serve as Co-Investigator on Earthcube Cyberconnector project which will make Unidata data available to a wide range of research and education communities beyond the traditional Unidata community.

Status Report: Python

April 2017 - October 2017

Ryan May, John Leeman, Sean Arms, Julien Chastang, Ward Fisher

Activities Since the Last Status Report

Python Training Efforts

- In lieu of the annual training workshop, we held regional workshops at Millersville University, University of Oklahoma, and University of Wisconsin. These were well attended, with a total of around 70 attendees.
- We are offering a short course at the 2018 AMS Annual Meeting: Python for Dynamical Meteorology using MetPy. This is an attempt to go beyond the traditional introductory Python course and demonstrate our tools with a practical example.
- Ryan May and John Leeman will be travelling to Texas Tech University to hold an intermediate Python training workshop in November 2017.
- We continue to improve and expand the Unidata Python Workshop, ensuring it stays current with the latest developments in the scientific Python stack. Teaching the regional workshops, as well as preparing for the Texas Tech workshop and anticipation of teaching the 2018 annual workshop, have revealed challenges in the size of the various teaching units, so additional revision to the materials will be done to ensure that the units are more easily exchanged and reduced. The materials also continue to be updated to expose the latest developments in MetPy and Siphon.
- John Leeman has spear-headed “MetPy Mondays”, as weekly series of posts and screencasts on the Unidata Developers’ blog. These highlight applying our Python collection of tools to various problems (or other related topics). The goal is to keep these to short chunks that can be digested in 5-10 minutes.

Progress has been made on the following:

- The re-worked training materials for the Python workshop worked well in practice.
- Expansion of available training materials and in-person training offerings
- MetPy Mondays have generated positive feedback as a way of engaging our community of users interested in Python

MetPy

The MetPy community continues to grow. There have been several more externally-driven Pull Requests, both for bug fixes and new features. The MetPy [twitter account](#) has reached 330 followers, and we are seeing mentions and use of MetPy on social media with increasing regularity. MetPy had two releases over the last few months. 0.5.1 fixed a few bugs in 0.5 (in advance of the last two regional workshops), while 0.6 added many new features (such as isentropic interpolation and many severe weather indices). Many of these features were created by our two summer interns, Tyler Wixtrom and Matt Wilson. That release also had additions from four external contributors. We also moved towards hosting multiple versions of the documentation on GitHub Pages, including a separate version for development. This is especially important as we have had support requests from users trying to use unreleased

features with released versions of the software.

We were successful in both MetPy-related proposals to NSF submitted last spring. The first is a collaboration between Ryan May and John Leeman of Unidata with Kevin Goebbert of Valparaiso University, funded by the Office of Advanced Cyber-infrastructure's Software Infrastructure for Sustained Innovation program. The focus of this proposal is to advance MetPy's functionality to make it a true replacement for GEMPAK. The second proposal, funded through EarthCube, also involves Ryan May, and is a collaboration between Unidata and CISL within UCAR, as well as Lamont-Doherty Earth Observatory at Columbia University and Anaconda Inc. (formerly Continuum Analytics). This work focuses on integrating several Python tools (Xarray, Dask, and MetPy) to enable Python-powered analytics of climate data on platforms like the Cheyenne supercomputer.

We continue to try to enhance the openness of MetPy's development model. Ideas for further development are captured in the GitHub issue tracker, and we are expanding our use of GitHub milestones for planning releases, and we welcome feedback. We hope to use the milestones as a way to communicate our plans to the community, and have the corresponding issues used for feedback. General foci over the next few releases will be dominated by the work needed for the NSF awards, including: xarray integration (likely including hooks into xarray for unit support), expanding calculations, more data formats, and simplified, declarative-style, plotting capabilities. Additional developments will also be driven by needs identified in the process of updating training materials.

Progress has been made on the following:

- Multiple versions of MetPy documentation are available, specifically separating development and release documentation
- Moving MetPy's support for upper air data to Siphon to simplify MetPy's testing needs, eliminating remote data access
- Community contribution: fixing equivalent potential temperature, updating METAR code mapping for weather symbols, and adding sigma and frontogenesis calculations
- Two MetPy-related NSF awards
- Community awareness and involvement progressing well in year 3

Siphon

We have begun to expand the scope of Siphon beyond only supporting the TDS. This has started by moving the support for downloading upper air datasets out of MetPy. In addition to simplifying MetPy's testing, it gives us a proper location to host code for accessing non-TDS-hosted datasets. This is important as we continue to refine our Python training offerings to be widely applicable and represent more real-world use cases.

We have also made some simple improvements to Siphon's API based on the patterns of access we see in the training materials. While it is still a bit involved to access TDS catalogs from Python, these improvements have simplified the process and we continue to look for ways to simplify data access.

Progress has been made on the following:

- Expanding Siphon's scope beyond TDS to a variety of useful atmospheric science

datasets.

- Enhancing Siphon's API to simplify the process of accessing data from TDS catalogs.

External Participation

The Python team attends conferences as well as participates in other projects within the scientific Python ecosystem. This allows us to stay informed and to be able to advocate for our community, as well as keep our community updated on developments. We continue to engage with the [Pangeo](#) project, which is a grass-roots effort to develop a community stack of tools serving the atmospheric, oceanic, land, and climate science communities. This engagement is enhanced by the start of work on the EarthCube collaborative award.

As participants in a broader Open Source software ecosystem, the Python team regularly encounters issues in other projects relevant to our community's needs. As such, we routinely engage these projects to address challenges and submit fixes. Ryan May continues to be an active participant in the Python community, serving as a core member of the matplotlib development team and contributing features and fixes to cartopy, xarray, and pint, among others. These contributions have solved many problems encountered by our community or by Unidata staff while working on other projects. We also continue to host Jeff Whittaker's netCDF4-python project repository; Jeff continues to be the active maintainer of the project.

Progress has been made on the following:

- Fixed multiple problems in CartoPy related to the projection of GOES-16 data
- Fixed a variety of issues in matplotlib, including some related to wind barbs as well as unit-handling
- Contributed code to integrate pint and matplotlib to simplify plotting with units
- Added support to xarray to enhance its interoperability with Siphon

Ongoing Activities

We plan to continue the following activities:

- Unidata Python training workshop
- Growing Siphon as a tool for accessing remote data sets, including those hosted on THREDDS data servers
- Growing and developing MetPy as a community resource for Python in meteorology
- Continued participation in the scientific Python community as advocates for the atmospheric science community
- Working with JupyterHub as a way to facilitate data-proximate analysis
- MetPy Mondays for engaging the community
- As resources and time permit, continue growing the Online Python Training project by writing Jupyter notebooks specifically targeted towards teaching the geoscience community programming concepts.

New Activities

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

- Present MetPy and Siphon work and roadmap at AMS 2018
- Teach short course on MetPy at AMS 2018

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

- Attend SciPy 2018
- Offer Python as part of the regular Unidata training workshop (the first time after restructuring for the regional workshops)

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

- Evaluate the possibility of extending siphon functionality to interface with the AWIPS-II EDEX server
- Restructure our annual Python workshop to be a full week with introduction to Python/git, intermediate with MetPy/Siphon/etc., and developer hack-day

Areas for Committee Feedback

We are requesting your feedback on the following topics:

1. Are there any additions you'd like to make to MetPy's or Siphon's development roadmap?
2. What are the biggest obstacles that you see to the use of Python with other Unidata technologies, or for use in meteorology in general?
3. We continue to maintain the Unidata Python Workshop with fresh, relevant, and up-to-date content. However, we would welcome feedback from our committees on topics we may not be covering in the workshop.

Relevant Metrics

Siphon

- 97% test coverage
- Watchers: 15
- Since 1 April 2017:
 - Active Issues: 35 (18 created, 16 closed)
 - Active PRs: 14 (14 created, 14 closed)
 - External Issue Activity: 2 opened, 2 comments
 - External PR Activity: 2 opened, 0 comments
 - Unique external contributors: 6
 - Stars: 10 (51 total)
 - Forks: 0 (17 total)
 - Commits: 61
- Since 1 October 2016
 - Active Issues: 41 (26 created, 19 closed)
 - Active PRs: 25 (23 created, 23 closed)
 - External Issue Activity: 4 opened, 5 comments
 - External PR Activity: 2 opened, 0 comments
 - Unique external contributors: 8

- Stars: 25 (51 total)
- Forks: 0 (17 total)
- Commits: 122

MetPy

- 97% test coverage
- Watchers: 32
- Since 1 April 2017
 - Active Issues: 125 (86 created, 57 closed)
 - Active PRs: 128 (105 created, 103 closed)
 - External Issue Activity: 19 opened, 52 comments
 - External PR Activity: 11 opened, 19 comments
 - Unique external contributors: 25
 - Stars: 66 (198 total)
 - Forks: 2 (81 total)
 - Commits: 308
- Since 1 October 2016
 - Active Issues: 218 (160 created, 123 closed)
 - Active PRs: 195 (187 created, 184 closed)
 - External Issue Activity: 29 opened, 89 comments
 - External PR Activity: 18 opened, 46 comments
 - Unique external contributors: 35
 - Stars: 119 (198 total)
 - Forks: 5 (81 total)
 - Commits: 630

Unidata Python Workshop

- Watchers: 19
- Since 1 April 2017
 - Active Issues: 64 (43 created, 40 closed)
 - Active PRs: 56 (52 created, 54 closed)
 - External Issue Activity: 0 opened, 2 comments
 - External PR Activity: 0 opened, 0 comments
 - Unique external contributors: 1
 - Stars: 12 (63 total)
 - Forks: 0 (51 total)
 - Commits: 133
- Since 1 October 2016
 - Active Issues: 82 (62 created, 53 closed)
 - Active PRs: 71 (71 created, 70 closed)
 - External Issue Activity: 2 opened, 2 comments
 - External PR Activity: 0 opened, 0 comments
 - Unique external contributors: 3
 - Stars: 26 (63 total)
 - Forks: 1 (51 total)
 - Commits: 269

Unidata Online Python Training

- Watchers: 4
- Since 1 April 2017
 - Active Issues: 0 (0 created, 0 closed)
 - Active PRs: 0 (0 created, 0 closed)
 - External Issue Activity: 0 opened, 0 comments
 - External PR Activity: 0 opened, 0 comments
 - Unique external contributors: 0
 - Stars: 6 (12 total)
 - Forks: 0 (2 total)
 - Commits: 1
- Since 1 October 2016
 - Active Issues: 9 (7 created, 8 closed)
 - Active PRs: 16 (16 created, 16 closed)
 - External Issue Activity: 0 opened, 0 comments
 - External PR Activity: 0 opened, 0 comments
 - Unique external contributors: 0
 - Stars: 12 (12 total)
 - Forks: 0 (2 total)
 - Commits: 17

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
 Python can facilitate data-proximate computations and analyses through Jupyter Notebook technology. Jupyter Notebook web servers can be co-located to the data source for analysis and visualization through web browsers. This capability, in turn, reduces the amount of data that must travel across computing networks.
2. **Develop and provide open-source tools for effective use of geoscience data**
 Our current and forthcoming efforts in the Python arena will facilitate analysis of geoscience data. This goal will be achieved by continuing to develop Python APIs tailored to Unidata technologies. Starting with the summer 2013 Unidata training workshop, we developed an API to facilitate data access from a THREDDS data server. This effort has been encapsulated with the new [siphon](#) project, which is an API for accessing remote data, including the THREDDS data server. Moreover, Python technology coupled with the HTML5 Jupyter Notebook technology has the potential to address "very large datasets" problems. Jupyter Notebooks can be co-located to the data source and accessed via a web browser thereby allowing geoscience professionals to analyze data where the data reside without having to move large amounts of information across networks. This concept fits nicely with the "Unidata in the cloud" vision and the goals outlined [Unidata 2018 Five-year plan](#). Lastly, as a general purpose programming language, Python has the capability to analyze and visualize diverse data in one environment through numerous, well-maintained open-source APIs. The additional development of [MetPy](#) fills the need for domain-specific analysis and visualization tools in Python.
3. **Provide cyberinfrastructure leadership in data discovery, access, and use**
 The TDS catalog crawling capabilities found in siphon will facilitate access to data remotely served by the Unidata TDS, as well as other TDS instances around the world.

4. **Build, support, and advocate for the diverse geoscience community**

Based on interest from the geoscience community, Unidata, as part of its annual training workshop, now hosts a three day session to explore [Python with Unidata technology](#). Also, to advance the use of NetCDF in Python, Unidata has promoted Jeff Whitaker's [NetCDF4-python project](#), including hosting its repository under Unidata's GitHub account. Unidata is initiating a project to provide [online Python training](#) specifically targeting geoscience students. Unidata is also fostering some community development of meteorology-specific tools under the MetPy project.

Prepared *September 2017*

Status Report: Support

April 2017 - October 2017

Jennifer Oxelson, Tom Yoksas, UPC Staff

Activities Since the Last Status Report

Training

- In 2017 The UPC will be focusing its in-person training efforts on regional workshops
- Additional resources will be directed towards developing online training materials.

New Activities

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

- Enhance electronic support offerings
- Create instructional materials for online virtual training

Relevant Metrics

Since January 26, 2006 over 54060 user support "transactions" (new inquiries and follow-ups) have been processed through the Unidata inquiry tracking system. Other methods of providing answers to questions posed (e.g., github, Stack Overflow, Jira, mailing list replies, etc.) add substantially to the support load.

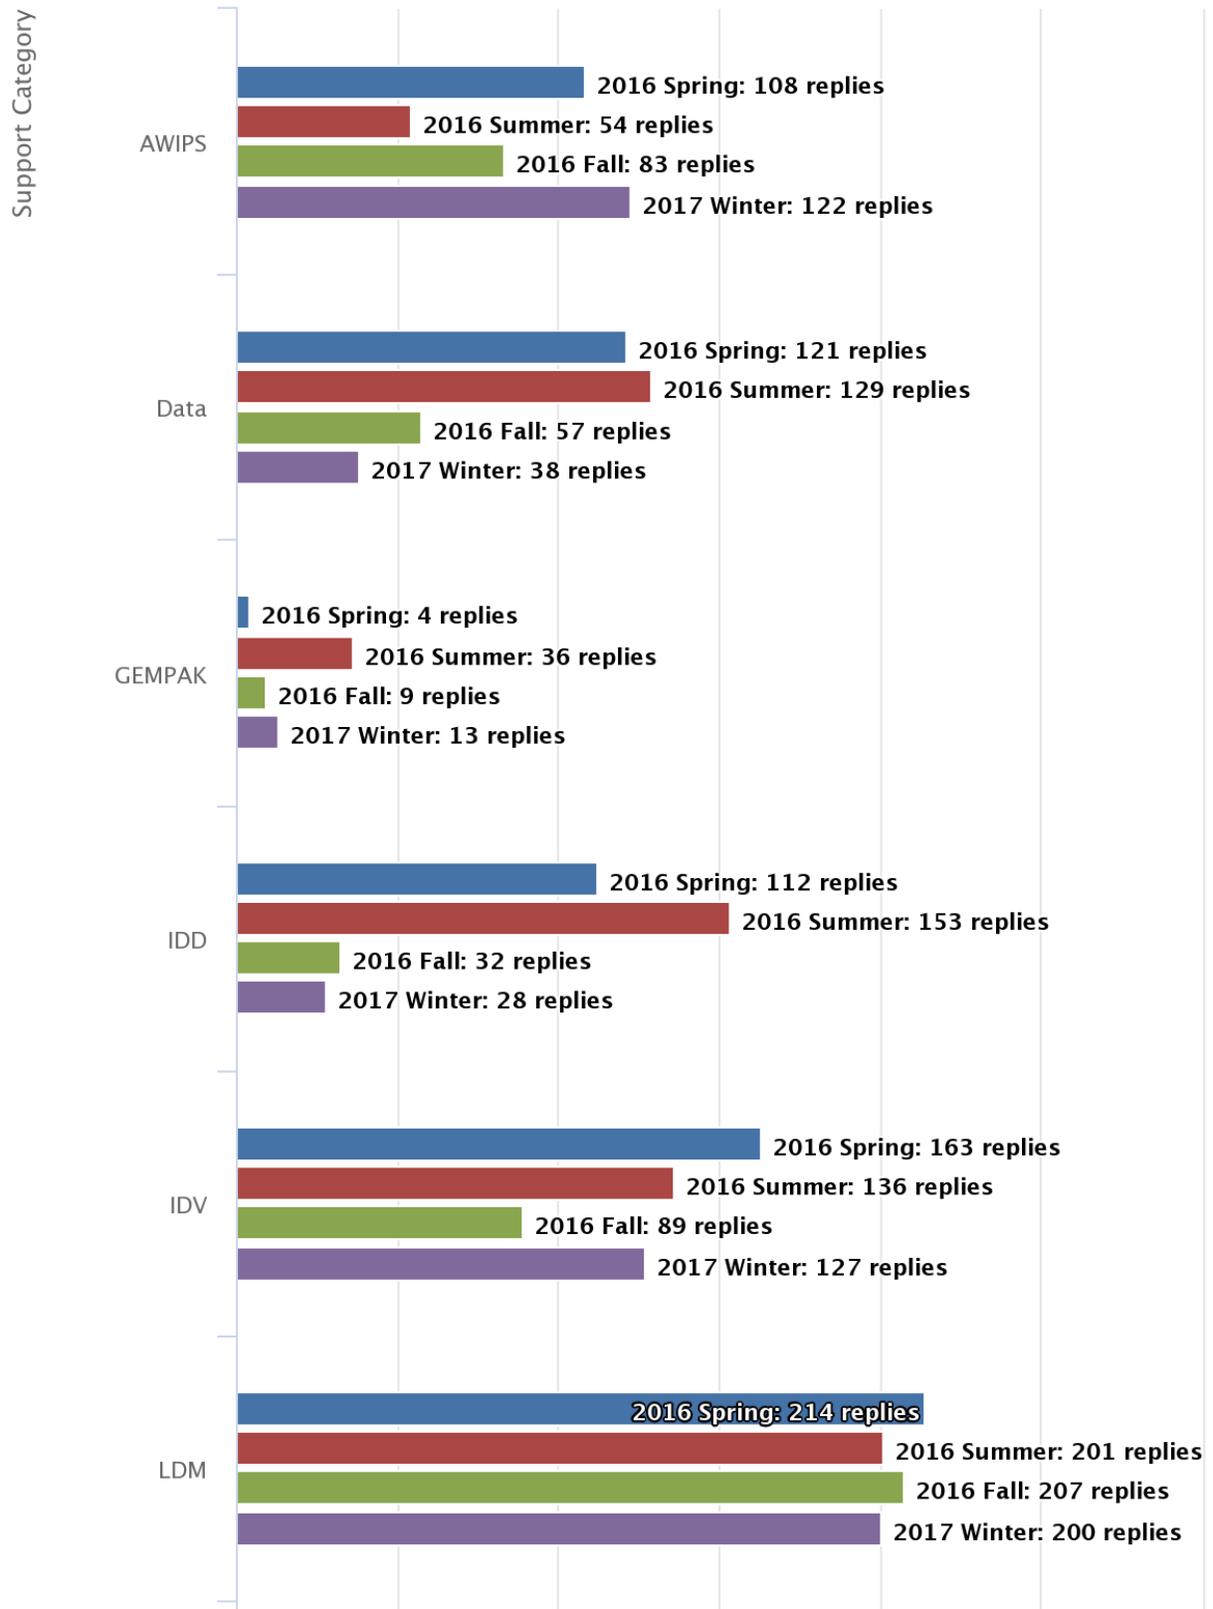
Additional metrics may be found in the [Comprehensive Metrics Data](#) portion of this meeting's agenda.

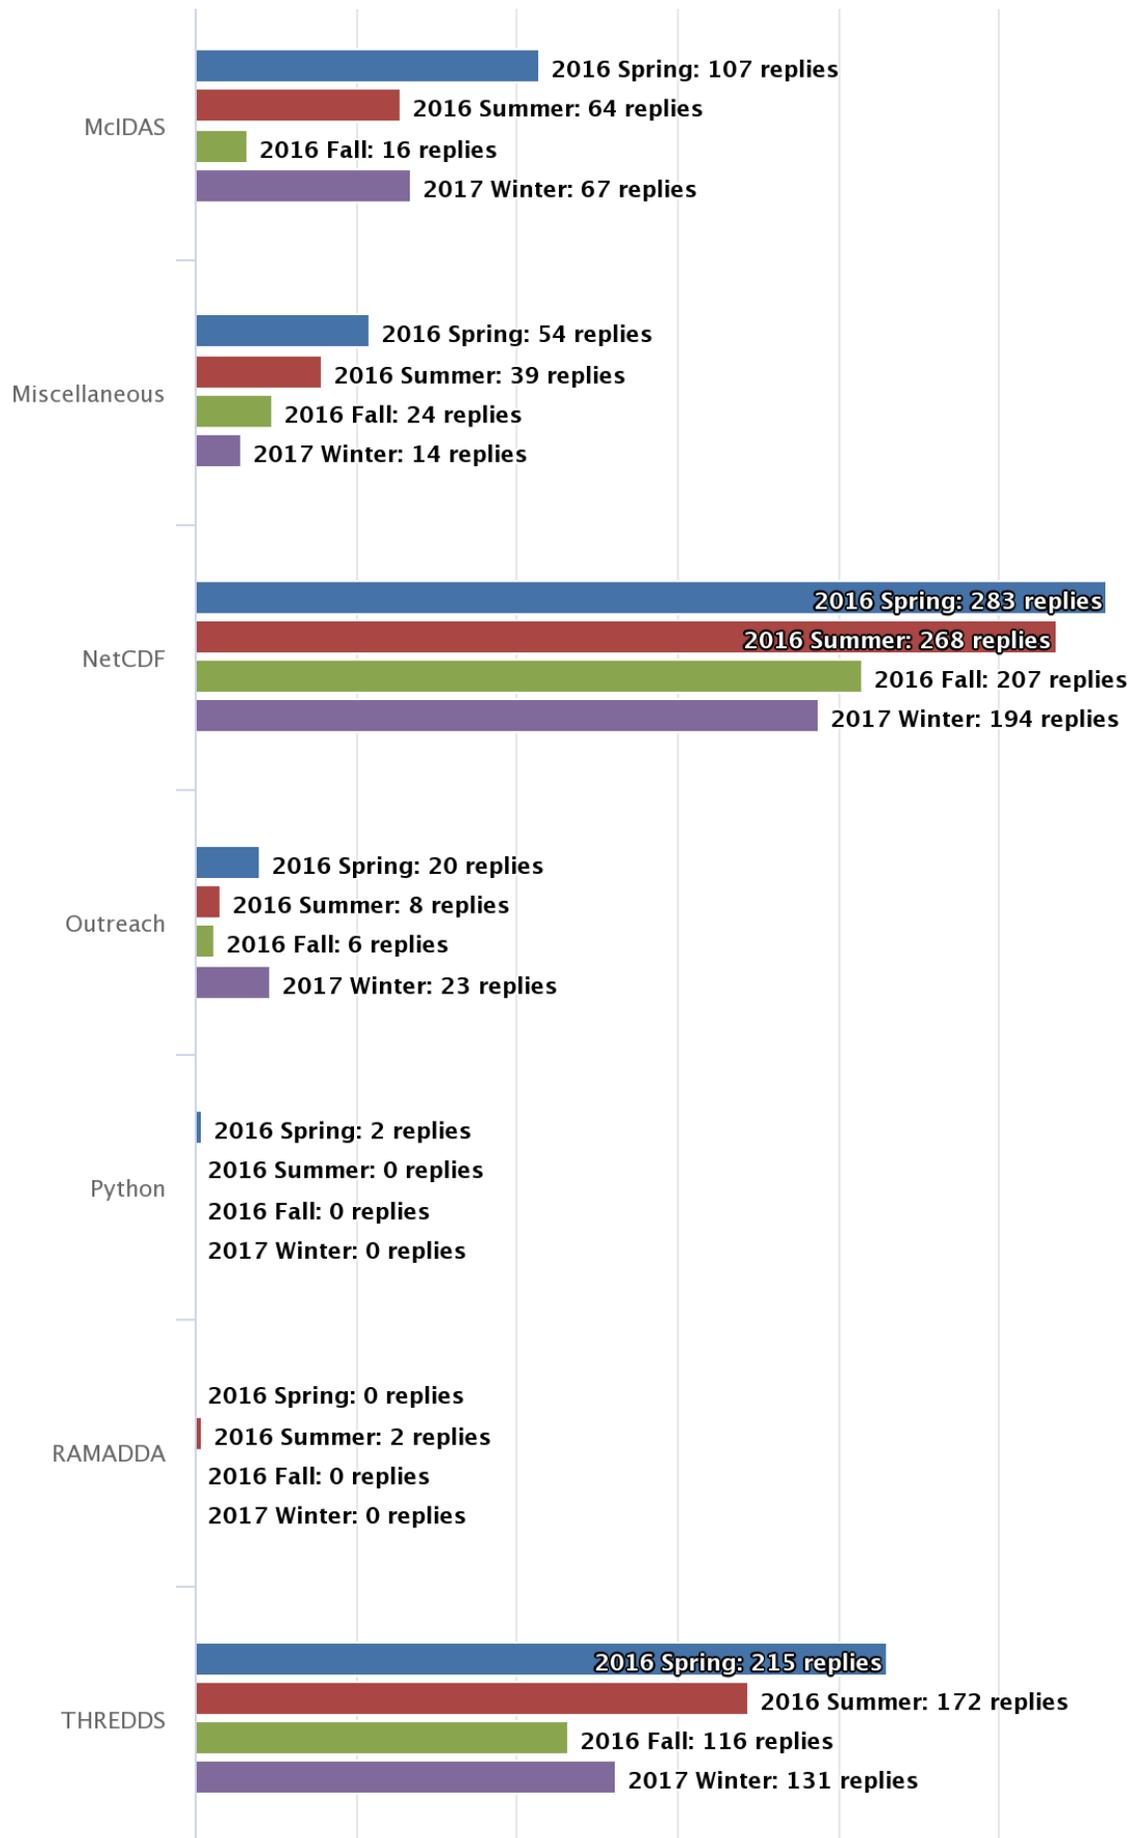
Fig. 1: Below are histograms that portray the number of Unidata email responses for categories of support logged in the Unidata Inquiry Tracking System for the 15 month period from October 1, 2016 until September 30, 2017.

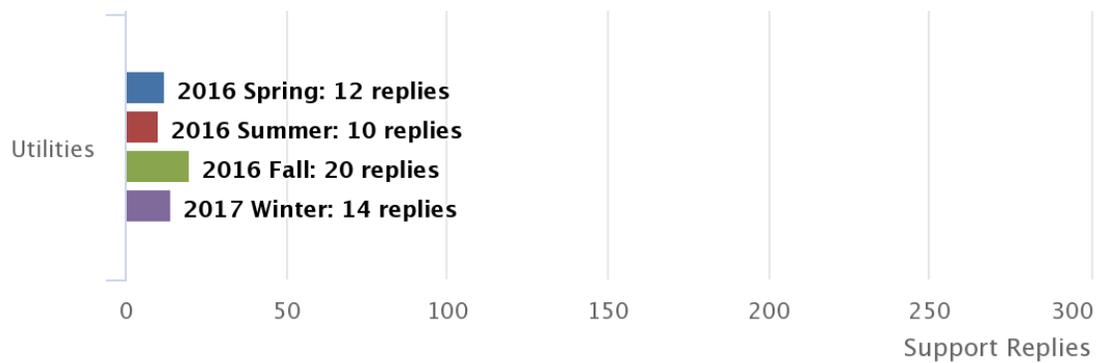
The quarters shown are defined as: Winter: January, February, March; Spring: April, May, June; Summer: July, August, September; Fall: October, November, December.

Total Number of Support Replies by Support Category per Quarter

April 1, 2016 to April 1, 2017







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

Category	eSupport Departments
AWIPS	Support AWIPS
Data	Support CaseStudy, Support CONDUIT, Support Datastream, Support LEAD, Support Level II, Support NOAAPORT, Support SUOMINET
GEMPAK	Support GEMPAK
IDD	Support IDD, Support IDD Antarctica, Support IDD Brasil, Support IDD Cluster, Support IDD SCOOP, Support IDD TIGGE
IDV	Support IDV, Support IDV Storm, Support McV, Support VisAD
LDM	Support LDM
McIDAS	Support McDevelop, Support McIDAS
Miscellaneous	Administration, Development, Plaza, Staging Folder, Support, Support eSupport, Support Miscellaneous, Support Platforms, Support Plaza, Systems
NetCDF	Support LibCF, Support netCDF
Outreach	Outreach, Polcomm, Support Egrants, Support News, Support Outreach, Support Workshop, Usercomm

Python	Support Python
RAMADDA	Support RAMADDA
THREDDS	Support netCDF Java, Support THREDDS
Utilities	Support LDM-McIDAS, Support netCDF Decoders, Support netCDF Perl, Support OPeNDAP, Support Rosetta, Support UDUNITS

Comments

- The total support provided by the UPC remains substantial, and yearly totals have been relatively constant for the past several years. Overall support activities vary by somewhat by quarter. Spikes in support for individual packages is largely correlated with the release of new distributions.
- Support for netCDF continues to be substantial, and is understandable given the very large number of users of the package worldwide.
- Support for the legacy visualization packages GEMPAK and McIDAS has decreased over the past year, probably due to GEMPAK users investigation use of AWIPS.
- Support for AWIPS has been increasing steadily and now exceeds that for GEMPAK.
- Support for Python scripting using MetPy is growing markedly.
- Support for LDM, IDD, and Data continue at a high levels and show some variability throughout the year.
- The high numbers for outreach reflect the high level of activity in a variety of activities including organizing sessions at various national meetings.

Notes

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

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

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
Unidata User Support enables access to geoscience data by supporting the use of tools created and/or supported by the UPC.
2. **Build, support, and advocate for the diverse geoscience community**
The user support provided by the UPC is recognized throughout the atmospheric science community. Unidata's outreach efforts are routinely called out in surveys of the NCAR/UCAR community.

Prepared *October 2, 2017*

Status Report: THREDDS

April 2017 - October 2017

Sean Arms, Ethan Davis, Dennis Heimbigner, Ryan May, Christian Ward-Garrison

Activities Since the Last Status Report

The THREDDS Project

The THREDDS Project encompasses four projects: netCDF-Java, the THREDDS Data Server (TDS), Rosetta, and Siphon (the Unidata Python client to interact with a TDS). For specific information on Siphon, please see the Python Status Report. An update regarding cloud efforts related to the TDS can be found in the Cloud Computing Activities Status Report.

Released netCDF-Java / TDS version 4.6.10 (Stable)

Progress has been made on the following:

- The 4.6.x line of development is now in maintenance mode so that the team can focus on v5.0. However, “maintenance mode” is taking up quite a bit of resources and progress on v5.0 has been impacted.

Focus netCDF-Java / TDS (Soon-to-be Beta) v5

Our last update indicated that the THREDDS team was preparing to release a beta version of the THREDDS Data Server (version 5.0) at the end of May. Unfortunately, due to external projects coming to a close, we were unable to meet that deadline. We hope to have the beta out before the beginning of 2018...*real soon now*TM

Progress has been made on the following:

- The Nexus Repository Manager at <https://artifacts.unidata.ucar.edu> has been upgraded from version 2 to version 3 and it will now host **all** build artifacts. For users, this means:
 - Developers should update their software builds to pull netCDF-Java/TDS artifacts from the Maven repository at <https://artifacts.unidata.ucar.edu/repository/unidata-all/>.
 - Documentation for versions 5 and later of netCDF-Java/TDS will reside at <https://artifacts.unidata.ucar.edu/repository/thredds-doc/> and—more conveniently—at <http://thredds-doc.unidata.ucar.edu/>.
- The configuration management tool [Ansible](#) has shown great promise as a way for users to be able to deploy TDS and other Unidata software in an automated fashion.
- DAP4 in the TDS has been updated to be consistent with the specification and to successfully allow the netCDF-C DAP4 and NetCDF-java libraries to read DAP4 responses from the TDS.
- New Coverage data type allows for subsetting across array boundaries (often called

the “seam” problem)

- Uses the new edal-java based ncWMS 2.0 server, as well as javascript client Godiva3
- CatalogScan feature allows for incremental updating of TDS catalogs without the need to restart Tomcat
- Upload/Download support has been added to TDS. This now includes an upload web form accessible as <http://.../thredds/upload>.
- Unit and Integration tests are passing in 5.0. This is a big step towards releasing a beta.
- ncSOS has been integrated into the TDS distribution (as part of the OIIP project - see the Rosetta section for more details)
- The access to the netCDF-c library via jni is now thread-safe so that the HDF5 library no longer needs to be built with thread-safe support.

Dependencies, challenges, problems, and risks include:

- The longer the 4.6.x line of development is maintained, the longer it will take to move forward with the 5.x line of development

Rosetta

Rosetta is progressing thanks to support from a NASA ACCESS grant (the Oceanographic In-situ data Interoperability Project, or **OIIP**), in which Unidata is partnering with the PO.DAAC at JPL and UMASS-Boston.

Progress has been made on the following:

- Support for the NCEI NODC netCDF v2.0 templates (metadata standards)
- Extension of the NCEI templates to support metadata critical to the use of electronic tagging datasets
- Support automated transformation of output from electronic animal tagging datasets in the Electronic Tag Unified File Format (eTUFF) format via Rosetta.
- Working to create a unified workflow for the gui wizard interface that allows for selection of which metadata standards to use when determining recommended/required metadata
- Engaging with the netCDF Linked Data initiative to define best practices identifying netCDF metadata to a particular metadata standard.

Dependencies, challenges, problems, and risks include:

- Two of the core javascript libraries used by Rosetta have been abandoned by their original creators. One has been picked up by the community (SlickGrid), while the other is in limbo (jWizzard). Unidata will likely need to pick up jWizzard and maintain it for use within Rosetta, at least internally. However, it would be a good community service to open this up to a wider audience, but resources would be required to do so.

Ongoing Activities

We plan to continue the following activities:

- Documentation updates - We are reworking the tutorial material for the TDS v5.0 with the goal of enabling asynchronous training. The material will undergo a major overhaul to include the use Docker containers, video snippets, and other new forms of training tools.
- Maintain thredds.ucar.edu and keep up with the addition of new datasets to the IDD
 - GOES-16 data, with tiles stitched together using python, available on our test TDS.
- Continue development of the TDS python client siphon, as well as extending its functionality to interface with other web services and servers

The following active proposals directly involve THREDDS work:

- The NASA ACCESS award with JPL is entering into the second year of the two year award. The award is titled: "Leveraging available Technologies for Improved interoperability and visualization of Remote Sensing and in-situ Oceanographic Data at the PO.DAAC" and was submitted with JPL/PO.DAAC. [Rosetta]
- EarthCube award: "Advancing netCDF-CF for the Geosciences". This two-year, Unidata lead project will work to extend netCDF-CF conventions in ways that will broaden the range of earth science domains whose data can be represented.
- Finished the second and final year of EarthCube award: "CyberConnector: Bridging the Earth Observations and Earth Science Modeling for Supporting Model Validation, Verification, and Intercomparison" with George Mason University.

New Activities

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

- Releasing a beta of TDS 5.0
- Officially advertising a public TDS 5.0 Test Server [currently found at <http://thredds-test.unidata.ucar.edu/thredds/catalog.html>]
- Engaging with Axiom Data Science to obtain datasets for use in prototyping a GRID featureCollection.

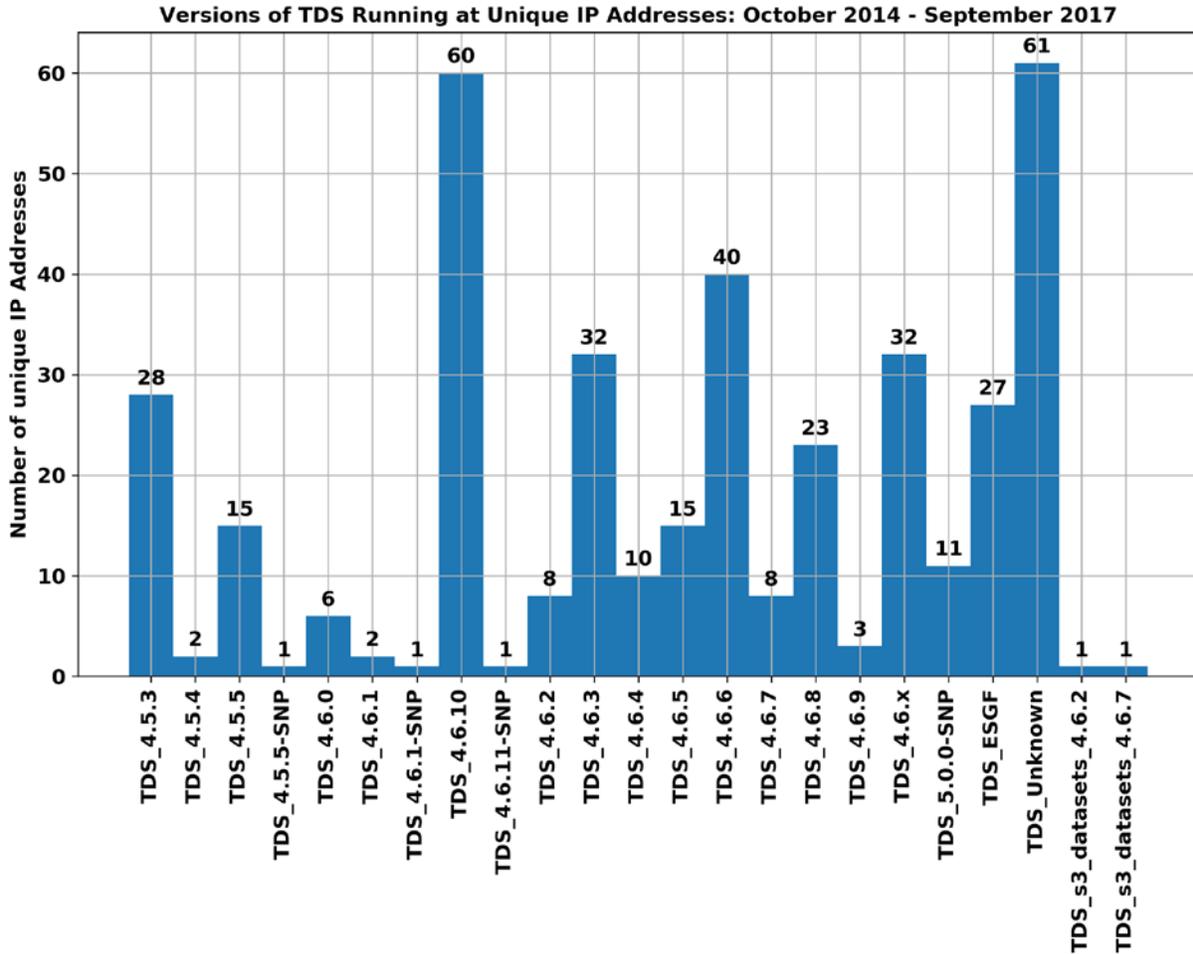
Over the next twelve months, we plan to organize or take part in the following:

- Support the [UGRID spec](#) within the CDM
- Support the [netCDF-CF simple-geometry spec](#) within the CDM
- Create a TDS plugin layer for external services
 - Finalize the TDS plugin layer.
 - Upgrade the ncWMS, ncISO, and other plugin services to use the new TDS 5.x plugin layer
 - Incorporate ncSoS into TDS
- Transitioning thredds.ucar.edu to TDS 5.x
- Getting TDS v5.0 to a stable release
- Getting netCDF-Java v5.x to a stable release

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

- Enable Rosetta to publish to a TDS

Relevant Metrics



9558 unique IPs started up thredds from November 2014 through September 2017, **536** of which are publicly accessible servers. Publicly accessible is defined as the following URL patterns being accessible with an HTTP HEAD requests with a return status less than 400:

http://<ip address>/thredds/catalog.xml
 http://<ip address>:8080/thredds/catalog.xml

This information is only known for servers running v4.5.3 and above. There are many reasons why these number are so different. The differences could be due to:

- Reporting TDS running behind a firewall that does not allow incoming traffic on 80 or 8080 (the ports tested)
- It might be possible that a TDS running through a proxy server may not been “seen” in this analysis as publicly reachable at the normal url pattern (<server>/thredds/catalog.xml)

- A TDS running in the past is no longer running today
- Finally, the most likely reason: people testing the TDS on their local machine, but not actually running a server

Note that the vast majority of the publicly accessible servers are running v4.6.3 or above (v4.6.10 was the most current release during this period, and was released on 20 April 2017, and is the most commonly run version of the 4.6.x line of the TDS). This indicates that users and organizations running the TDS tend to follow along closely with the current releases of the TDS.

As with the last report, the updated analysis also indicates a number of sites are running TDS v5.0, even though it is pre-beta. This underscores the desire for the new features in 5.0, and highlights one reason why we feel the need focus most of our efforts on and to move all new development to the v5 codebase.

Note that there are some odd looking versions of the TDS being reported in the log files, such as TDS_4.28.x. It is likely these version numbers are actually generated by software that is being built on top of the TDS. Previous version of the figure above listed each of these odd versions as its own entry; these oddities are aggregated together and shown as "TDS_Unknown".

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**

The work of the THREDDS group is comprised of two main areas: the THREDDS Data Server (TDS) and the Common Data Model (CDM) / netCDF-Java library. The TDS provides catalog and data access services for scientific data using OPeNDAP, OGC WCS and WMS, HTTP, and other remote data access protocols. The CDM provides data access through the netCDF-Java API to a variety of data formats (e.g., netCDF, HDF, GRIB). Layered above the basic data access, the CDM uses the metadata contained in datasets to provide a higher-level interface to geoscience specific features of datasets, in particular, providing geolocation and data subsetting in coordinate space. The CDM also provides the foundations for all the services made available through the TDS.

The data available from the IDD is a driving force on both the TDS and netCDF-Java development. The ability to read all the IDD data through the netCDF-Java library allows the TDS to serve that data and provide services on/for that data.

2. **Develop and provide open-source tools for effective use of geoscience data**

Unidata's Integrated Data Viewer (IDV) depends on the netCDF-java library for access to local data, and on the THREDDS Data Server (TDS) for remote access to IDD data. At the same time, the CDM depends on the IDV to validate and test CDM software. Many other tools build on the CDM / netCDF-Java library (eg ERDDAP, Panoply, VERDI, etc) and on the TDS (ESGF, LAS, ncWMS, MyOcean, etc).

3. **Provide cyberinfrastructure leadership in data discovery, access, and use**

The Common Data Model (CDM) / netCDF-Java library is one of the few general-purpose implementations of the CF (Climate and Forecast) metadata standards.

Current active efforts in CF that we are involved with include use of the extended netCDF-4 data model (CF 2.0) and for point data (Discrete Sampling Geometry CF-DSG).

The TDS has pioneered the integration of Open Geospatial Consortium (OGC) protocols into the earth science communities. Strong international collaborations have resulted in WCS and WMS services as part of the TDS.

The CDM and TDS are widely used implementations of the OPeNDAP DAP2 data access protocol. Unidata has worked with the OPeNDAP group to design, develop, and implement a new version of the DAP specification, DAP4, which is now available in the TDS server and the netCDF-Java client software stack.

4. **Build, support, and advocate for the diverse geoscience community**

The THREDDS project is involved in several international standardization efforts (CF, OGC, etc.) which cross-cut a multitude of disciplines, both inside and outside of the geoscience community. The netCDF-Java client library, as well as the TDS often serve as incubators for new pushes in these efforts.

Prepared *October 2017*