Users Committee Spring Meeting
19-20 April, 2017
Agenda

Wednesday, 19 April 2017

8:30 - 9:00 Continental Breakfast

9:00 - 9:15 Administrative Items

- Introduction: New Members
- Date for Fall meeting
- Review of Action Items from September 2016 meeting
- Changes to UCAR travel system

9:15 - 9:30 Strategic Advisory Committee Report - Mohan Ramamurthy and Russ Schumacher

9:30 - 10:45 Director's Report & Questions - Mohan Ramamurthy

10:45 - 11:00 Break

11:00 - 11:30 Available GTS data - Josh Young & Becky Cosgrove

11:30 - 12:45 Lunch

12:45 - 1:30 Around-the-table reports from members

1:30 - 2:15 Status Reports - (Users Committee should review status reports prior to the meeting) - staff will be available for questions

2:15 - 2:30 Break

2:30 - 3:00 Cloud Updates - Jeff Weber

3:00 - 3:45 ERDDAP & Cataloguing across THREDDS - Rich Signell

3:45 - 4:30 GOES-16 - Tom

4:30 - 5:00 2018 Triennial Workshop - Josh

5:00 Adjourn

6:30 Collaborative discussion on the day's proceedings over dinner

- Dinner at West Flanders Brewery, 1125 Pearl Street Mall (map)

Thursday, 20 April 2017

8:30 - 9:00 Continental Breakfast

9:00 - 9:30 Strategic Plan Update - Mohan Ramamurthy

9:30 - 10:00 Site Visit Team Update - Mohan Ramamurthy

10:00 - 10:30 DeSouza Nominees - Josh

10:30 - 10:45 Break

10:45 - 11:45 Blue Sky Session

Status Report: AWIPS and GEMPAK

October 2016 - April 2017

Michael James

AWIPS Activities Since the Last Status Report

AWIPS in the Cloud

Unidata continues to provide access to real-time AWIPS data in the cloud, using first an Azure for Research grant, and then the Jetstream grant (outlined in the Cloud Activities status report) once the Azure grant ran its course. RedHat 7 development has also been migrated from Azure to Jetstream, and a 16.4.1 el7 release is in development as of the writing of this status report.

OS X Support

Unidata has supported CAVE for OS X since 15.1.1, though support for all viz plugins and perspectives is still in progress, with recent updates making available the National Centers Perspective (NCP) on macOS using Java 1.8 and Eclipse 4.5.

Three obstacles are currently being addressed to make the Mac client more functional:

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).
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.
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).

AWIPS for Windows

Attempts were made to release a stable Windows CAVE 16.2.2 client, but the efforts were shelved due to too much time being wasted. When the 16.4.1 source code is ready for Linux and Mac, an attempt will be made to again export a Windows client from the Eclipse product. Users are unhappy with this delay. The Unidata AWIPS developmer also is unhappy with this delay.

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 16.4.1 repositories, and a Python 3-compliant version has been made available by PyPI (installable via pip) and integrated into the Python Training Workshop.

Data Delivery / OGC Plugin Investigation

Experimentation with the Data Delivery and OGC plugins has been shelved until more time or resources are available.

Unidata AWIPS Release Notes

16.4.1 TBD (April-May 2017)

- The manual ingest endpoint has been moved from /awips2/edex/data/manual to /awips2/data_store/ingest in order to accomodate a distributed LDM-EDEX environment, where UNIWISC McIDAS area files and FNEXRAD gini composites should be written to a shared directory (/awips2/edex/data/manual is unmountable across Azure Red Hat 7 systems)
- New RPM builds for python (2.7.11), java (1.8), Eclipse (4.5.1), pupynere (1.0.15), werkzeug (0.12.1), cherrypy (3.8.2), tpg (3.2.2), cython (0.25.2), metpy (0.5.0), ldm (6.13.6), groovy (2.4.9), shapely (1.5.17), six (1.10.0), postgres (9.3.16), gdal (2.1.3).
- GFE suite server site backup functionality has been completely removed.
- All geospatial data filters have been removed (which were filtering out southern hemisphere records)
- D2D Procedures have returned in 16.4.1 (open,save,delete,new)
- GOES-16/R data ingest via the NIMAGE feed is now supported.
- GOES-16/R level2 products are still unavailable to Unidata, but I have a contact at NWS who will make available Python utlities for converting from level1b to level2.
- GOES-16/R and Himawari menus have been removed from the awips2-goesr repo and integrated into com.raytheon.viz.satellite.
- Hydroapps and MPE processes have either been muted or removed from install completely.

Unidata GEMPAK 7.3.1 (November 2016)

- Released as source code tarball, 64-bit Linux RPM, and 64-bit RPM compiled against Python 2.7 (gempak-python).
- Incorporated NAWIPS 7.3.1 table, map, and code updates.
- Released a gzipped tarball for macOS Darwin (gempak-osx-7.3.1.tar.gz)

Relevant Metrics

Software Downloads

- AWIPS Downloads by Month, Jan-Apr 2017
- GEMPAK Downloads by Month, Jan-Apr 2017
- AWIPS Downloads by Affiliation
- AWIPS Downloads US vs. non-US
- AWIPS Downloads Core vs. non-Core
- AWIPS Downloads by site for Core Institutions, Jan-Apr 2017
- AWIPS Downloads by site for EDUs, Jan-Apr 2017
- more...

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 efficiently accessing geoscience data. The cloud-based data server maintained by Unidata invites a number of users and we are seeing more adoption with time, as the software becomes more stable, more platforms are supported, more data is served conveniently and freely to the community.*

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.

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 *April 2017*

Status Report: Cloud Computing Activities

October 2016 - April 2017
Sean Arms, Julien Chastang, Ethan Davis, Steve Emmerson, Ward Fisher, Michael James, Ryan May, Jennifer Oxelson, Mohan Ramamurthy, Mike Schmidt, Christian Ward-Garrison, Jeff Weber, Tom Yoksas

Unidata technical staff have deployed experimental and production software in several cloud computing environments. 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. 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. In addition, we have we have an experimental "motherlode" class server running in the NSF XSEDE Jetstream cloud serving a subset of the IDD data via a TDS and RAMADDA. These data are supplied by an LDM relay also running on the Jetstream cloud. (Note to readers, we are in the process of transitioning from our Start up to Research allocation on the Jetstream cloud. As such, it is possible the Jetstream links in this report may be down at the time of this reading. We expect to complete the transition in late Spring or early Summer.) Also on Jetstream, Unidata is experimenting with AWIPS EDEX server running in the cloud.

Activities Since the Last Status Report

Docker Development

With the goal of better serving our core community and in fulfillment of objectives articulated in *Unidata 2018: Transforming Geoscience through Innovative Data Services* , Unidata is investigating how its technologies can best take advantage of cloud computing. To this end, we have been employing Docker container technology to streamline building, deploying, and running Unidata technology offerings in cloud-based resources. 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, and the commercial clouds of Microsoft Azure and Amazon AWS. Our preliminary efforts are available on various Docker-related Unidata GitHub and DockerHub repositories, and cloud demonstration servers.

Progress has been made on the following:

- **LDM**: Continue to make minor improvements to the LDM container that is currently running in the Jetstream cloud as a relay node and leaf node and supplying data to RAMADDA and THREDDS Data Server containers.
- **RAMADDA**: We continue to make minor improvements to the RAMADDA Docker container. This container is running without issue on the Jetstream Cloud.
- **THREDDS Data Server and TDM**: we collaborated with Axiom Data Science to construct a TDS/TDM container that will satisfy Unidata and Axiom objectives and avoid duplication of effort. We have a demonstration server running on the Jetstream Cloud. We continue to make minor enhancements to this container as needed.
- **ADDE**: We have a demonstration server running on the Jetstream cloud for AMS 2017 meeting. We have informally asked UW/SSEC to allow us to make freely available a containerized distribution of the ADDE server component of McIDAS-X. A formal petitioning is still ongoing.
- **CloudIDV**: We are in the process of obtaining feedback from beta users. We are exploring new uses for CloudIDV beyond bringing the IDV to new classes of device. Currently we are using it for data-proximate visualization; this involves working with data sets so large that standard desktop and laptop machines do not have enough resources. This is offset by hosting the CloudIDV instance in the Cloud on Large (32+ GB Ram, 16+ multiple processors, etc) VMs.
- **CloudStream**: 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.
- **Python:** We have updated the Python Docker container as environment for running our tools on JupyterHub in the cloud. We are evaluating resources for running this machine, and plan to work on increasing the capabilities (e.g. dashboards) for this environment.

Dependencies, challenges, problems, and risks include:

While these efforts are promising initial steps, there are challenges ahead in making these technologies useful to our community. Apart from client technologies like the CloudIDV and Jupyter notebooks, it is unlikely that most of our users will initially 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. Moreover, we may have to rethink workflows in a cloud environment (data-proximate analysis and visualization, for example) in addition to porting present Unidata cyberinfrastructure to the cloud.

Unidata Cloud Grants, Awards and Resources

Microsoft Azure Awards

Microsoft awarded two $20,000 "Azure for Research Grants" to Unidata in 2016. While these grants are time-limited, they provide us with invaluable resources to experiment with cloud computing environments. We are successfully running the THREDDS Data Server, RAMADDA, CloudIDV and EDEX servers in the Microsoft Azure Cloud.

Progress has been made on the following:

- We have deployed numerous services and instances to the Azure Cloud, mirroring our experiments with the Amazon cloud infrastructure.
- While these Azure grants have been invaluable for testing and experimentation, we are transitioning to the Jetstream cloud to continue our progress in this area.

XSEDE Jetstream Award

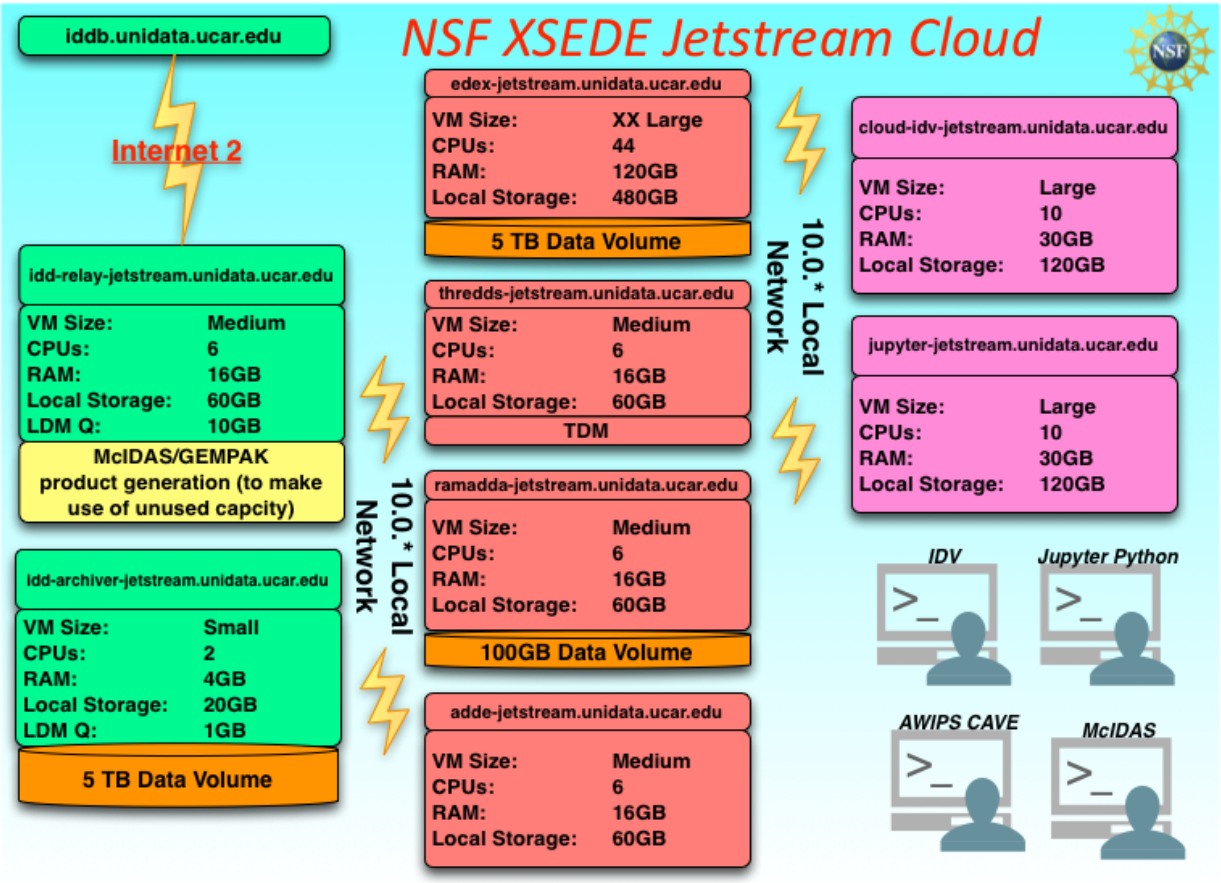

To further investigate how the Unidata community can benefit from Unidata technologies in the cloud, Unidata obtained a large XSEDE “Research” grant on the Jetstream

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. We wish to continue our research of porting Unidata technology into a variety of cloud environments including non-commercial, research-oriented clouds such as Jetstream. Specifically, we would like to deploy a motherlode class machine on the Jetstream cloud with Docker technology in a manner similar to what we accomplished with our Azure resources. Jetstream became available in February of 2016. We continue to experiment with Jetstream initially on our "Start up" grant which is now almost completely exhausted and our "Research" grant which is getting underway.

In March 2017 the public-facing EDEX server edex-cloud.unidata.ucar.edu was successfully migrated from the Azure cloud to Jetstream, with a significant performance improvement due to the larger instances available through Jetstream.

Progress has been made on the following:

We presented the progress made under the Start Up grant at the Seattle AMS annual meeting in January. We are in the process of transitioning from our Start Up to Research grant. We have a new github repository to capture this effort. We expect to have the plan detailed in the diagram above in place this summer for the ESIP Summer meeting in Bloomington, IN.

Dependencies, challenges, problems, and risks include:

The transition from Start Up to Research grant is going smoothly, and we would like to accelerate this transition to make maximum use of our XSEDE resources. We hope, this outcome will put us in a strong position to ask for additional resources when our grant period ends.

Amazon Awards

Progress has been made on the following:

- Learning about Amazon's cloud infrastructure
- Designing an initial architecture to support storing all NEXRAD-2 data in Amazon's cloud
- Started transferring GOES-16 GRB data to S3 bucket
- Began building configurations from TDSConfig on Travis and hosting in Amazon S3

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.

CloudIDV Application Streaming

Unidata has received a second year of Azure resources from Microsoft under the "Azure for Research" program. The primary focus of this award is continue work on creating an application-streaming platform for the IDV and other Unidata technologies. Secondary focus is on testing Unidata services in the Azure cloud, and examining the performance of Azure when hosting Docker instances.

Progress has been made on the following:

- We have created a Dockerized version of the IDV bundled with a remote desktop/application streaming server. We are currently finishing up the first version of the associated web dashboard, "CloudControl".
- We have released the new versions of the Dockerized IDV, CloudIDV.
- We have updated the generic application-streaming container for use by our community with their own legacy software, "CloudStream" to allow for encrypted-only and shared sessions.

AWIPS in the Cloud

The Azure for Research Grant for Unidata AWIPS has allowed for the edex-cloud open data server to live on, with an on-site EDEX server available as a replacement for those periods of time where a cloud-based server is not funded. This grant has enabled the development of a RedHat 7 supported EDEX and CAVE build, which can take advantage of the Azure file sharing architecture to create a distributed EDEX environment, scalable to data requirements.

A similar EDEX Data Server has been maintained by Unidata for Embry-Riddle Aeronautical University (ERAU) on an Amazon EC-2 instance, though access is restricted only to ERAU domains.

One of these Azure for Research grants was awarded in order to support a cloud-based EDEX data server for the UCAR community (which ran live through April 2017), and to help support cloud-based testing and development for a distributed EDEX server (summarized at AMS 2017: https://ams.confex.com/ams/97Annual/webprogram/Paper315787.html). This grant was essential to Unidata AWIPS development and support on 64-bit Red Hat/CentOS 7 systems. Though the Azure grant is now at the end of its life, the XSEDE Jetstream Award is

now used to continue cloud-based AWIPS development and serving of real-time data to the UCAR community.

Ongoing Activities

We plan to continue the following activities:

- Continue to create the LDM/IDD FNEXRAD (NEXRAD Level III national composite radar products) and UNIWISC (GOES-East/West image sectors) in a virtual machine hosted in AWS
- Use the LDM to move NEXRAD Level II data into AWS S3 buckets in real-time
- Develop TDS access to data stored in S3
- Maintain the TDS on AWS serving level II radar data
- Deploy Jupyter notebooks that provide access to NEXRAD Level II data stored in S3
- Continue experimenting with our Jetstream resources for running Unidata technology in the cloud and staging motherlode class machine.
- Migrate Red Hat/CentOS 7 development to Jetstream virtual machines.

Big Data Project

- 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.

- 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

New Activities

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

- Deploy new versions of CloudIDV and CloudStream
- Deploy a motherlode-class machine in the Jetstream cloud under our XSEDE Research Grant. Demonstrate this capability at the ESIP Summer meeting in Summer 2017.
- Work with the Amazon Big Data Project team to bring GFS model output into the cloud in real time.

- Begin work on new Amazon Web Services grant to install and maintain an EDEX server in a cloud environment for the academic and research community to access
- Continue containerizing Unidata software and advocating for its use through presentations and blog entries.
- Look into AWS grant to support efforts into jupyterhub research project

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

- Implement machine images of our software for easy deployment in a virtual environment.
- Make full use of our XSEDE Jetstream Research grant and eventually ask for a renewal of this grant for 2018.
- Investigate cloud-based streaming services for CAVE deployment.

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

While Unidata is successfully moving its technology offerings to the cloud, we have not reinvented our technology to best take advantage of cloud computing. We hope to better research this area in the long-term.

Areas for Committee Feedback

- What clouds are our community using, either commercial or private?
- What new cloud technologies are our community using/investigating on their own initiative?
- Who would like to volunteer to beta test CloudIDV?

Relevant Metrics

Docker image downloads are available from Unidata's Dockerhub repository. Especially popular are the THREDDS Docker container with 465 downloads and CloudIDV with 306 downloads.

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 and private cloud services will allow subscribers to those services to access data quickly and at low cost.

2. **Develop and provide open-source tools for effective use of geoscience data**
   Running existing Unidata-developed and supported tools and processes (e.g. IDV, EDEX, RAMADDA, generation of composite imagery) in a range of cloud environments makes these tools and data streams available to cloud service subscribers at low cost. It also gives us insight into how best to configure existing and new tools for most efficient use in these environments.

3. **Provide cyberinfrastructure leadership in data discovery, access, and use**
   Unidata is uniquely positioned in our community to experiment with provision of both data and services in the cloud environment. 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.

Appendix

- Cloud Computing Projects at Unidata Spreadsheet
- Cloud Computing Resources at Unidata Spreadsheet

Prepared *April, 2017*

Status Report: Community Services

October 2016 - April 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:

- Unidata's Jupyter Notebook Gallery
- Unidata Staff at AGU Fall 2016 Meeting
- Unidata Software Training Workshops: Tell Us What You Think
- Unidata Summer Student Internship Available!
- Call for Proposals: Unidata 2017 Community Equipment Awards
- Python-Focused Software Training Workshop at Millersville University
- Unidata Program Center Welcomes John Leeman
- AMS 2017 Conference Highlights from the Unidata Staff
- Wanted: Student Representative for Unidata Users Committee
- 2017 DeSouza Award Nominations
- Python-Focused Software Training Workshop at University of Oklahoma
- Check Out this Podcast with Unidata's MetPy Developers
- Unidata Communications Internship Available
- 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:

- Assisting the Strategic Advisory Committee in their efforts to provide guidance for the next Unidata Strategic Plan.
- Continued work on the pilot project to assist community members with data management and document the resulting workflows for the DMRC was awarded

supplemental funding by NSF. Five academic volunteers participated in the first phase (Florida Institute of Technology, Lyndon State College, North Carolina State University, Millersville University, and the University of Wyoming). As a part of this project, community services intern Larissa Gordon joined the UPC through September 2016. The second phase of this project is currently underway with the recruitment of an intern to work on evaluating additional freely-available open source tools for inclusion in the DMRC.

- Work in providing a pilot online training (http://unidata.github.io/online-python-training/) focused on Python-content. This was the secondary project for community services intern Larissa Gordon (along with other staff)
- Community engagement at recent professional society conferences (2016 AGU Fall Meeting and AMS Annual Meeting)
- 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
- Evaluate additional open source freely-available tools for inclusion in the DMRC and seek additional case studies with associated outreach through social media and the Unidata blog
- Support Committees work towards producing the next Strategic Plan

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:

- 51,121 unique pageviews (52,386 in previous period)

- 8.4% of total unique pageviews (11.8% in previous period)

Top community pages

1. All blog pages
   39592 unique pageviews (39443 in previous period)
   77% of total community pageviews (75% in previous period)
2. www.unidata.ucar.edu/events
   4045 unique pageviews (6113 in previous period)
   8% of total community pageviews (12% in previous period)
3. www.unidata.ucar.edu/community
   3814 unique pageviews (4119 in previous period)
   7% of total community pageviews (6% in previous period)
4. www.unidata.ucar.edu/about
   2971 unique pageviews (3316 in previous period)
   6% of total community pageviews (6% in previous period)

Social media statistics, March 31, 2017

1. # of Twitter followers: 631 (up from 552)
2. # of Facebook followers: 645 (up from 568)

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 *April 2017*

Status Report: Community Equipment Awards

Sponsored by the National Science Foundation

October 2016 - April 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 has been formed and will meet on April 20 at the Unidata Program Center to review and recommend which proposals to fund.

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 87 awards totaling close to $1,100,000.

Prepared March 2017

Status Report: GOES-R

October 2016 - April 2017

Mike Schmidt, Jeff Weber, Tom Yoksas

Activities Since the Last Status Report

This is a ***new*** status report.

Ongoing Activities

We plan to continue the following activities:

- Ingest the GOES ReBroadcast (GRB) from GOES-16 in real-time on the new 4.5 m satellite dish that we installed on the eastern most (UNAM) pad at the NCAR Mesa Lab using the NOAA funded UW/SSEC/CIMSS Community Satellite Processing Package (CSPP) for Geostationary Data (GEO) package
    - The UNAM pad was used since it can support satellite dishes of a much greater size than needed and because of existing conduit from the pad into the Mesa Lab that leads to a UCAR Network Engineering and Telecommunications Section (NETS) closet in which the receiver for the GRB (Quorum GRB-200) is housed.
    - CSPP GEO is responsible for GRB ingest and productization (e.g., stitching together of broadcast tiles)
- Copy all ABI imagery, Space Weather products, Geostationary Lightning Mapper (GLM) and Level 2 products to one or more S3 buckets 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
- Work with NCAR/RAL to provide GOES-16 data for their science objectives
- Reach out to other NCAR labs to promote use of GOES-16 data
- 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
- Create an LDM feed of ABI imagery for distribution within NCAR/UCAR
    - Due to its sheer volume, we are ***not*** planning on distributing GOES-16 imagery, Geostationary Lightning Mapper (GLM), space weather or Level 2 products in the IDD

- Investigate 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
- 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 revamped using GOES-16 data
- Pursue funding to install GOES-S ingest and data serving capabilities

Relevant Metrics

- Lots O Data!

  We are currently receiving and saving on the order of 500 GB per day and a variety of products including Geostationary Lightning Mapper (GLM) and Level 2 products are not yet being broadcast!

Questions for Committee Members

- What new coverages would should be added to the UNIWISC feed?
- What kind(s) of data access methods are desired?
- 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 have been bundled by Raytheon 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 *April 5, 2017*

Status Report: Internet Data Distribution

October 2016 - April 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  65350.716 M bytes/hour
Average hourly volume  41292.195 M bytes/hour

Average products per hour     413387 prods/hour

Feed                            Average             Maximum     Products
                       (M byte/hour)            (M byte/hour)   number/hour
FSL2                    11225.988    [ 27.187%]    17572.843    17696.844
CONDUIT                  8565.003    [ 20.742%]    22337.741    97060.289
NGRID                    6670.759    [ 16.155%]    12496.926    41718.889
NEXRAD2                  6103.271    [ 14.781%]     7829.229    71278.133
NOTHER                   4358.971    [ 10.556%]     7991.873     9203.044
NEXRAD3                  2108.416    [  5.106%]     2655.176    98827.200
FNMOC                    1298.238    [  3.144%]     4152.749     3616.667
HDS                       304.961    [  0.739%]      463.071    26127.956
NIMAGE                    181.822    [  0.440%]      374.231      196.356
GEM                       179.070    [  0.434%]      761.532     1135.000
FNEXRAD                   118.286    [  0.286%]      141.539      105.156
IDS|DDPLUS                 72.025    [  0.174%]       86.217    45798.556
UNIWISC                    68.762    [  0.167%]      115.008       48.444
EXP                        34.889    [  0.084%]       65.395      276.289
LIGHTNING                   1.735    [  0.004%]        3.327      297.778
```

Ongoing Activities

We plan to continue the following activities:

- Unidata continues to receive 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 (3) of university sites on **hrrr.unidata.ucar.edu**. Since HRRR and ESTOFS data were added to the NOAAPort Satellite Broadcast Network (SBN) in late September, 2014, continuing to relay the HRRR ingested from NOAA/GSD is considered to be of lesser importance and will be discontinued if the sites receiving the NOAA/GSD data lose interest.

- The HRRR is being experimentally served at:
  http://thredds-jumbo.unidata.ucar.edu/thredds/modelsHrrr.html
  (.xml for machine access)

- Other data sets we continue to explore with NOAA/GSD/ESRL are:
  - FIM
  - HIWPP

- HRRR fields and forecasts times that are not included in the NOAAPort expansion will be evaluated as additions to the CONDUIT IDD datastream.

- The UPC continues to relay FNMOC and CMC data model output directly to the community. FNMOC provides the COAMPS and NAVGEM model output and the CMC provides the GEM model output. Unidata has provided access to these data for the past 8 years, but on a "point-to-point" basis. GEM model output was converted from GRIB1 to GRIB2 in January, 2015. The CMC is now feeding output of their new hi-resolution (15 km) GEM model to Unidata who, in turn, relays the data to IDD participants.

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 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 **585** 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 relay node, **idd.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.

  Data input to the cluster nodes averages around 29 GB/hr (~0.7 TB/day). Over the period from January 12 through March 31, 2017 the average volume of LDM/IDD data flowing from the UCAR/NCAR network exceeded 2.9 Gbps (~31 TB/day), and peak rates reached 5.9 Gbps (which would be ~63 TB/day if the rate was sustained).

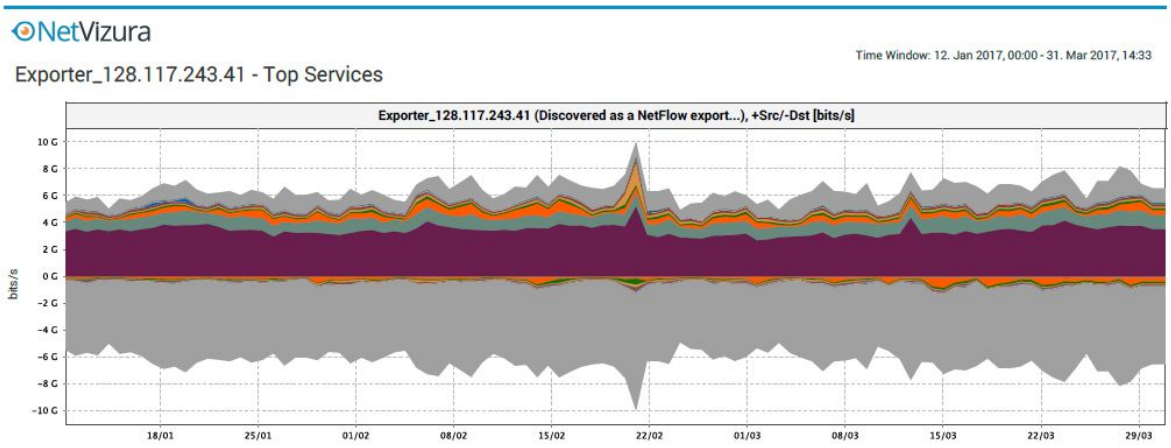

Services	Src		Dst		Total	
	Avg	Max	Avg	Max	Avg	Max
LDM [388]	2.9 Gbps (52.9 %)	5.2 Gbps	54.3 Mbps (1.0 %)	84.8 Mbps	2.9 Gbps (26.9 %)	5.3 Gbps
Unknown [80]	743.4 Mbps (13.7 %)	1.3 Gbps	10.4 Mbps (0.2 %)	40.5 Mbps	753.7 Mbps (6.9 %)	1.3 Gbps
HTTPS [443]	182.3 Mbps (3.4 %)	840.5 Mbps	71.1 Mbps (1.3 %)	735.9 Mbps	253.4 Mbps (2.3 %)	1.1 Gbps
SSH [22]	97.5 Mbps (1.8 %)	279.7 Mbps	66.1 Mbps (1.2 %)	464.4 Mbps	163.7 Mbps (1.5 %)	565.9 Mbps
Unknown [2049]	98.7 Mbps (1.8 %)	1.6 Gbps	34.8 Mbps (0.6 %)	194.9 Mbps	133.5 Mbps (1.2 %)	1.7 Gbps
Unknown [873]	36.3 Mbps (0.7 %)	122.8 Mbps	6.4 Mbps (0.1 %)	135.8 Mbps	42.7 Mbps (0.4 %)	210.9 Mbps
Unknown [2101]	37.4 Mbps (0.7 %)	50.0 Mbps	1.7 Mbps (<0.1 %)	11.5 Mbps	39.1 Mbps (0.4 %)	52.1 Mbps
Unknown [112]	17.6 Mbps (0.3 %)	35.6 Mbps	300.4 kbps (<0.1 %)	878.6 kbps	17.9 Mbps (0.2 %)	36.3 Mbps
Unknown [516]	196.8 kbps (<0.1 %)	1.7 Mbps	17.6 Mbps (0.3 %)	208.2 Mbps	17.8 Mbps (0.2 %)	209.2 Mbps
Unknown [53]	11.3 Mbps (0.2 %)	22.4 Mbps	5.3 Mbps (<0.1 %)	6.3 Mbps	16.7 Mbps (0.2 %)	28.4 Mbps
Unknown [5440]	15.9 Mbps (0.3 %)	286.6 Mbps	546.3 kbps (<0.1 %)	9.1 Mbps	16.5 Mbps (0.2 %)	288.5 Mbps
Unknown [5445]	102.7 kbps (<0.1 %)	6.5 Mbps	15.1 Mbps (0.3 %)	34.8 Mbps	15.2 Mbps (0.1 %)	34.9 Mbps

Cluster real server backends and accumulator nodes routinely have instantaneous output volumes that exceed a Gpbs. 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 increase in IDD data volume over the past year is largely attributable to the 0.25 degree GFS data in CONDUIT, the overall increase in the volume of data being transmitted in NOAAPort (which now routinely exceeds 13.6 GB/hr), and the increase in dual polarization NEXRAD data.

Questions for Committee Members

- Suggestions regarding content of data streams like CONDUIT, FNEXRAD, UNIWISC and 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.

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 being bundled by Raytheon in AWIPS-II.

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 experimenting with relaying data received in the IDD to Colombia.

Status Report: IDV with RAMADDA

October 2016 - April 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 (5.3) is the 4.6.7 (April, 2016). This was principally a bug fix release.

__ISL Changes__

- New api of ISL bundle attributes: bbox and time
- New `loadcatalog` tag to load a remote catalog URL
- New `publish` tag to publish a file to RAMADDA

__Java Version__

The nightly build version of the IDV is being 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 nightly build version of the IDV is being 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.

__Plugins__

Updated the location of the NOAA ESRL PSD plugin now that NOAA has shifted to HTTPS.

__Derived Quantities__

Added two APIs to calculate the vertical velocity based on hydrostatic and based on vertical pressure velocity

IDV Display Changes

__Latest Version of VisAD__

In the last six months, there have been several feature enhancements and bug fixes in VisAD:

- Improvements related to trajectory displays (ongoing project)

__3D trajectories Changes__

- Redo the 3D trajectory displays to improve the rendering performance and result (ongoing project)

__Satellite Imagery__

- Add new histogram feature to image display

__UI Changes__

- Allow entries for Alias, ParamGroups, ParamDefaults UIs to be sorted.
- Add view menu item to change the unit of altitude axis
- In the Profile Control, added a View menu item to change the unit of the altitude axis
- Create new Bundle menu item to dynamically link to the RAMADDA server

__Point Observations__

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

__WRF Hydro__

- Integrated the WRF-Hydro 1.1 version output with separated meta-data to generate a CF complaint point feature type dataset, and greatly improve the data loading and rendering performance in the IDV.

__NetCDF Point Cloud Display__

- Improvements for the handling of NetCDF point cloud display for WRF Hydro output

__Adaptive Resolution and Match Display Region__

- Add new APIs to dynamically relocate the bundle in the ISL script.

__IDV Web Cams__

- The IDV webcam list continues to be updated regularly

IDV WRF-Hydro Collaboration

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

IDV WRF-Tutorial Workshop

In the July WRF tutorial workshop, Yuan gave a 45 minutes presentation of the IDV main features and its applications in the WRF output datasets. We also help out the students in the classroom.

IDV Release

The last IDV release was 5.3 in May of 2016 and plan to release 5.4 in May 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.

IDV Publication Highlights

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

RAMADDA

Docker is a new cloud-centric technology that borrows from the notion of containers from the shipping industry to facilitate installation and deployment of server side applications. We have implemented a Docker container for easy distribution and installation of RAMADDA in a cloud environment. We coupled this effort with a Dockerized LDM with the goal of serving data in a cloud environment. We continue to work on this effort started last year and presented an update at the 2017 AMS annual meeting in Seattle, Washington.

IDV and RAMADDA Training, Conference Attendance and Presence

__2017 American Meteorological Conference (AMS) Annual Meeting__

- Presentation: WRF Hydro Stream Flow Display in the IDV
- Presentation: Geoscientific Data Distribution in the XSEDE Jetstream Cloud (RAMADDA)

__2016 American Geophysical Union (AGU) Annual Meeting__

- Poster: The Unidata Integrated Data Viewer
- Presentation: Data-proximate Visualization via Unidata Cloud Technologies

Ongoing Activities

We plan to continue the following activities:

__IDV in the Cloud__

With the goal of better serving our core community and in fulfillment of objectives articulated in “Unidata 2018: Transforming Geoscience through Innovative Data Services”, the IDV team will continue to investigate how its technologies can best take advantage of cloud computing. To this end, we have been employing Docker container technology to streamline building, deploying, and running Unidata technology offerings in cloud- based resources. Specifically, we have created Docker images for the IDV, RAMADDA, we are working on the LDM which, coupled with RAMADDA, will allow for the serving of real-time data in a cloud environment for IDV users. We have been experimenting with these Docker containers in the NSF XSEDE Jetstream Cloud

__IDV Instructional Videos__

We plan to continue producing more instructional videos on the IDV. 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.

New Activities

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

__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.

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 ~70 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 11,500 times compared with 10,000 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

October 2016 - April 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

- Data 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. It is expected that programmatic access to GOES-16 data served by ADDE will dwarf the access we have seen for the current GOES-GVAR imagery.
- Use of Unidata tools, especially netCDF, the IDV and GEMPAK, continues to grow internationally.
- IDD-Brazil continues to deliver data via the LDM in Africa.

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 *April 2, 2016*

Status Report: LDM

October 2016 - April 2017

Steve Emmerson, Tom Yoksas, Mike Schmidt, Julien Chastang

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:

- Improved HTML documentation based on latest workshop. Added graphics, troubleshooting section, and improved tips.
- Improved rotation of NOAAPort ingest logs.
- Improved multicast protocol (not yet ready).
- Improved logging of signals received by LDM server.
- Corrected handling of blank but non-empty lines in pqact(1) configuration-file.
- Merged changes to noaaportIngester(1) by Raytheon.
- Improved log message by LDM library when shmget(2) fails.
- Ported LDM package to FreeBSD.

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 proposal has been submitted to the NSF for this project. The program manager has approved it and we are waiting for approval by higher authorities.

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:

- Renew development and deployment of LDM-7 (assuming the proposal is funded)

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

- Development and deployment of LDM-7 (assuming the new proposal is funded)

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

October 2016 - April 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 add indexing to ADDE datasets to speed access into large datasets.

Current Activities

- Unidata McIDAS version 2017 is being prepared for release in May. v2017 includes all SSEC versions up to and including the current release, v2017.1 and Unidata updates and bugfixes.

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

The latest releases features the following:

- Updated ADDE servers for Himawari and GOES-R imagery and a preliminary release of a GLM (Geostationary Lightning Mapper) server
- ADDE server for GOES-R imagery: This server does __not__ work with the tiled GOES-16 image sectors that are being distributed in the NOAAPort SBN. A new ADDE server is being developed to support the NOAAPort-delivered images. It is hoped that the development can proceed quickly enough to be available in the Unidata v2017 release. If it is not, v2017 will be released early, and the new server will be made available in an addendum.

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 (I2) bandwidth usage by the McIDAS ADDE protocol routinely exceeds 8 TB/week.
- 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 release addressed building on newer OS versions. A new version will be readied

Geostationary Satellite Data Ingest

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 4 TB/month.
- Direct, programmatic access to real-time GOES-West (GOES-15) data via McIDAS ADDE routinely exceeds 2.3 TB/month.

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.

Continue investigating the feasibility of moving the GOES-R imagery and products to “the cloud” in near real-time. Discussions with Amazon Web Services representatives have already take place, and they are very interested in the GOES-R data being made available in the same way as Landsat imagery.

Strategic Focus Areas

We support the following goals described in Unidata Strategic Plan:

1. **Enable widespread, efficient access to geoscience data**
   *McIDAS remains **the** application of choice for the satellite meteorology community.*
   The Advanced Data Distribution Environment (ADDE) component of McIDAS 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 McIDAS, McIDAS-V, unlike its predecessors, is a fully open source application that is is in wide scale and growing use in the satellite meteorology community.
   McIDAS 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 serving over 4.5 TB per month. 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**
   McIDAS is sought for use by those interested in satellite meteorology worldwide.

Status Report: netCDF

October 2016 - April 2017

Ward Fisher, Dennis Heimbigner

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 50 open issues for netCDF-C, 10 open issues for netCDF-Fortran, and 12 open issues for netCDF-C++. The netCDF Java interface is maintained by the Unidata CDM/TDS group (which also uses Jira and GitHub), 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.

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.
- The hdf5 1.10 version generated, by default, backwards-incompatible binary netCDF4 files. This was addressable but was a short-notice high-priority issue which required immediate attention.

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.

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.

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?**
4. **Considering other modern code/software practices, in what area(s) do you feel netCDF is the most deficient? What avenue of modernization would be most practical for you?**

Relevant Metrics

There are currently about 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 **0.32** six months ago to **0.36** 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**.

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

Currently, Google hits, for comparison, are:

- **1,010,000** for netCDF-3
- **1,060,000** for netCDF-4
- **637,000** for HDF5
- **65,000** for GRIB2

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

- **263** for netCDF-3
- **475** for netCDF-4
- **8270** for HDF5
- **693** 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 *March 2017*

Status Report: Outreach Activities

October 2016 - April 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
- Represent Unidata and UCAR in OGC - 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

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

New Activities

No new activities are planned.

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.

Prepared *April 2017*

Status Report: Python

October 2016 - April 2017

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

Activities Since the Last Status Report

Python Training Efforts

- We continue to improve and expand the Unidata Python Workshop, ensuring it stays current with the latest developments in the scientific Python stack. Specifically, we are enhancing and adding Jupyter Notebooks in the geoscientific domain in addition to making sure we stay on top of infrastructure changes with the conda package manager, and the Jupyter environment.
- In lieu of the annual training workshop, we are holding several regional workshops in 2017. Millersville University, University of Oklahoma, and University of Wisconsin are confirmed host institutions.
- Ryan May (Unidata), Scott Collis (Argonne National Laboratory), Jonathan Helmus (Argonne National Laboratory), and Jed Sundwall (Amazon), taught a short course on using PyArt to access radar data from the S3 archive and combine that data with surface plots in MetPy
- We plan to submit a short course proposal to offer our a shortened version of our Python workshop at the 2018 AMS Annual Meeting
- Julien Chastang presented on Unidata's Online Python Training effort at the 2017 AMS Annual Meeting

Progress has been made on the following:

- Existing Python workshop materials have been significantly reworked to address shortcomings from an instructional point of view (and based on community feedback). Improvements include a more cohesive lesson plan, less duplication of material, and a more problem-driven approach to the instruction.
- Expansion of available training materials and in-person training offerings

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 220 followers, and we are seeing mentions and use of MetPy on social media with increasing regularity. MetPy had several bug fix releases (0.4.1, 0.4.2, 0.4.3), mainly fixing various issues with Skew-T support. Version 0.5 was released in advance of the first of the regional workshops; this release included bug fixes and more calculations (including some contributed ones), as well as our first API break (for LCL). Documentation has also been improved by using sphinx-gallery to generate the example gallery, reorganizing the API documentation, and moving to GitHub Pages (from ReadTheDocs). We have also recently enabled Google analytics on all our projects on GitHub pages, both for metrics and to understand what's being used.

We continue to improve MetPy's open 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 ZenHub to track issues status during our development sprints; both are open for any interested parties. 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 release or two include: reading point data (both BUFR and raw METAR) to facilitate generating NetCDF data for hosting on THREDDS, xarray integration (likely including hooks into xarray for unit support), and any work necessary to facilitate working with GOES-16 data. Our intention is for xarray integration to help reduce the amount of boilerplate code necessary for simple analysis scripts.

Progress has been made on the following:

- MetPy's examples have been ported to sphinx-gallery, which makes it much easier to build a gallery of examples, complete with descriptions, code highlighting, and links to API documentation
- MetPy's image tests have been improved to reduce false negatives
- Many fixes for Skew-T and Hodograph
- Migrated documentation from ReadTheDocs to GitHub Pages for sphinx-gallery support and to open opportunity for integration with Unidata branding
- Community awareness and involvement progressing well as we head into year 3
- Community contribution of virtual temperature, virtual potential temperature, and density calculations
- Two MetPy-related proposals submitted to NSF in March 2017

Siphon

Siphon represents our official Python support for TDS. While development has been slow of late, this is largely because its current capabilities meet current needs (versus some needs in MetPy). We anticipate developing Siphon to ensure that it is easy as possible to download data from a TDS in Python (such as crawling a server looking for data), keeping pace with new features added on the Java side. One specific development will be ensuring that Siphon works well with the 5.0 release of the TDS.

Progress has been made on the following:

- Support for interacting with TDS 5.0. Nothing really changed in Siphon, but bugs in the new version of TDS were revealed and fixed.
- Support added for reading THREDDS Client catalogs served by Hyrax servers
- Improved automated testing and build infrastructure to bring in line with that of MetPy.

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. Ryan May has also

continued to be an active participant in the matplotlib community, reviewing some pull requests and contributing several others. We also continue to host Jeff Whittaker's netCDF4-python project repository; Jeff continues to be the active maintainer of the project. Ryan May presented the Core Science Keynote for Python at the AMS annual meeting, as well as helped bring the future generation of students into fold by presenting on Python at the AMS Student Conference. We have also been engaging 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. Their goals for various software stack align well with ours, to the extent that MetPy was included in an Earthcube proposal submitted by this group.

Progress has been made on the following:

- Have continued to evaluate xarray (formerly xray, created by Stephen Hoyer) as a way to get CDM-like functionality in Python. It's current abilities provide a nice coordinate-aware data-object, as well as a way to attach attributes to arrays. The main hold up to integration with MetPy is being able to use with MetPy's unit support.
- Participated in the conda-forge project on GitHub; this is a community project to produce automated builds of conda packages using open recipes and infrastructure. We have contributed (and maintain) recipes for MetPy and Siphon, as well as their dependencies. These packages are available from the conda-forge channel on anaconda.org

Ongoing Activities

We plan to continue the following activities:

- Unidata Python training workshop
- Maintaining Siphon as an official Python API for working with TDS
- Growing and developing MetPy as a community resource for Python in meteorology
- Continued participation in the scientific Python community
- Relevant matplotlib support and fixes
- Working with JupyterHub as a way to facilitate data-proximate analysis
- Continue regular series of notebook-based blog posts on the Unidata Developer's blog to demonstrate the use of Python for various meteorological tasks
- As resources and time permits, continue making progress on 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:

- Attend SciPy 2017
- Hold multiple regional workshops focusing on Python

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

- 2018 AMS Annual Meeting in Austin, TX
    - Offer a version of our Python workshop as a short course
    - Python related presentations by Unidata staff
- Participate in another pyart/MetPy radar-focused short course

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

- 95% test coverage
- Watchers: 11
- Since 1 October 2016:
    - Active Issues: 10 (8 created, 3 closed)
    - Active PRs: 11 (9 created, 9 closed)
    - External Issue Activity: 2 opened, 3 comments
    - External PR Activity: 0 opened, 0 comments
    - Unique external contributors: 2
    - Stars: 16 (42 total)
    - Forks: 0 (15 total)
    - Commits: 64
- Since 1 April 2016
    - Active Issues: 25 (14 created, 6 closed)
    - Active PRs: 12 (16 created, 17 closed)
    - External Issue Activity: 2 opened, 3 comments

        - External PR Activity: 1 opened, 1 comments
        - Unique external contributors: 3
        - Stars: 22 (42 total)
        - Forks: 0 (15 total)
        - Commits: 96

MetPy

- 97% test coverage
- Watchers: 27
- Since 1 October 2016
    - Active Issues: 131 (74 created, 69 closed)
    - Active PRs: 90 (82 created, 85 closed)
    - External Issue Activity: 10 opened, 37 comments
    - External PR Activity: 7 opened, 28 comments
    - Unique external contributors: 15
    - Stars: 61 (142 total)
    - Forks: 0 (62 total)
    - Commits: 341
- Since 1 April 2016
    - Active Issues: 153 (124 created, 91 closed)
    - Active PRs: 127 (124 created, 122 closed)
    - External Issue Activity: 27 opened, 66 comments
    - External PR Activity: 24 opened, 47 comments
    - Unique external contributors: 33
    - Stars: 84 (142 total)
    - Forks: 14 (62 total)
    - Commits: 477

Unidata Python Workshop

- Watchers: 21
- Since 1 October 2016
    - Active Issues: 35 (19 created, 13 closed)
    - Active PRs: 21 (21 created, 18 closed)
    - External Issue Activity: 2 opened, 0 comments
    - External PR Activity: 0 opened, 0 comments
    - Unique external contributors: 2
    - Stars: 16 (53 total)
    - Forks: 0 (46 total)
    - Commits: 137
- Since 1 April 2016
    - Active Issues: 40 (33 created, 17 closed)
    - Active PRs: 38 (38 created, 35 closed)
    - External Issue Activity: 3 opened, 3 comments
    - External PR Activity: 0 opened, 0 comments
    - Unique external contributors: 3
    - Stars: 30 (53 total)

        - Forks: 1 (46 total)
        - Commits: 206

Unidata Online Python Training

- Watchers: 3
- 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
    - Forks: 0 (1 total)
    - Stars: 6 (6 total)
    - Commits: 16
- Since 1 April 2016
    - Active Issues: 40 (29 created, 35 closed)
    - Active PRs: 61 (61 created, 61 closed)
    - External Issue Activity: 1 opened, 1 comments
    - External PR Activity: 0 opened, 0 comments
    - Unique external contributors: 2
    - Stars: 6 (6 total)
    - Forks: 0 (1 total)
    - Commits: 94

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 communicating with a THREDDS 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 *March 2017*

Status Report: Support

October 2016 - April 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 (confirmed hosts: Millersville University, University of Oklahoma, and University of Wisconsin).
- 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 52120 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 January 1, 2016 until April 1, 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

January 1, 2016 to April 1, 2017

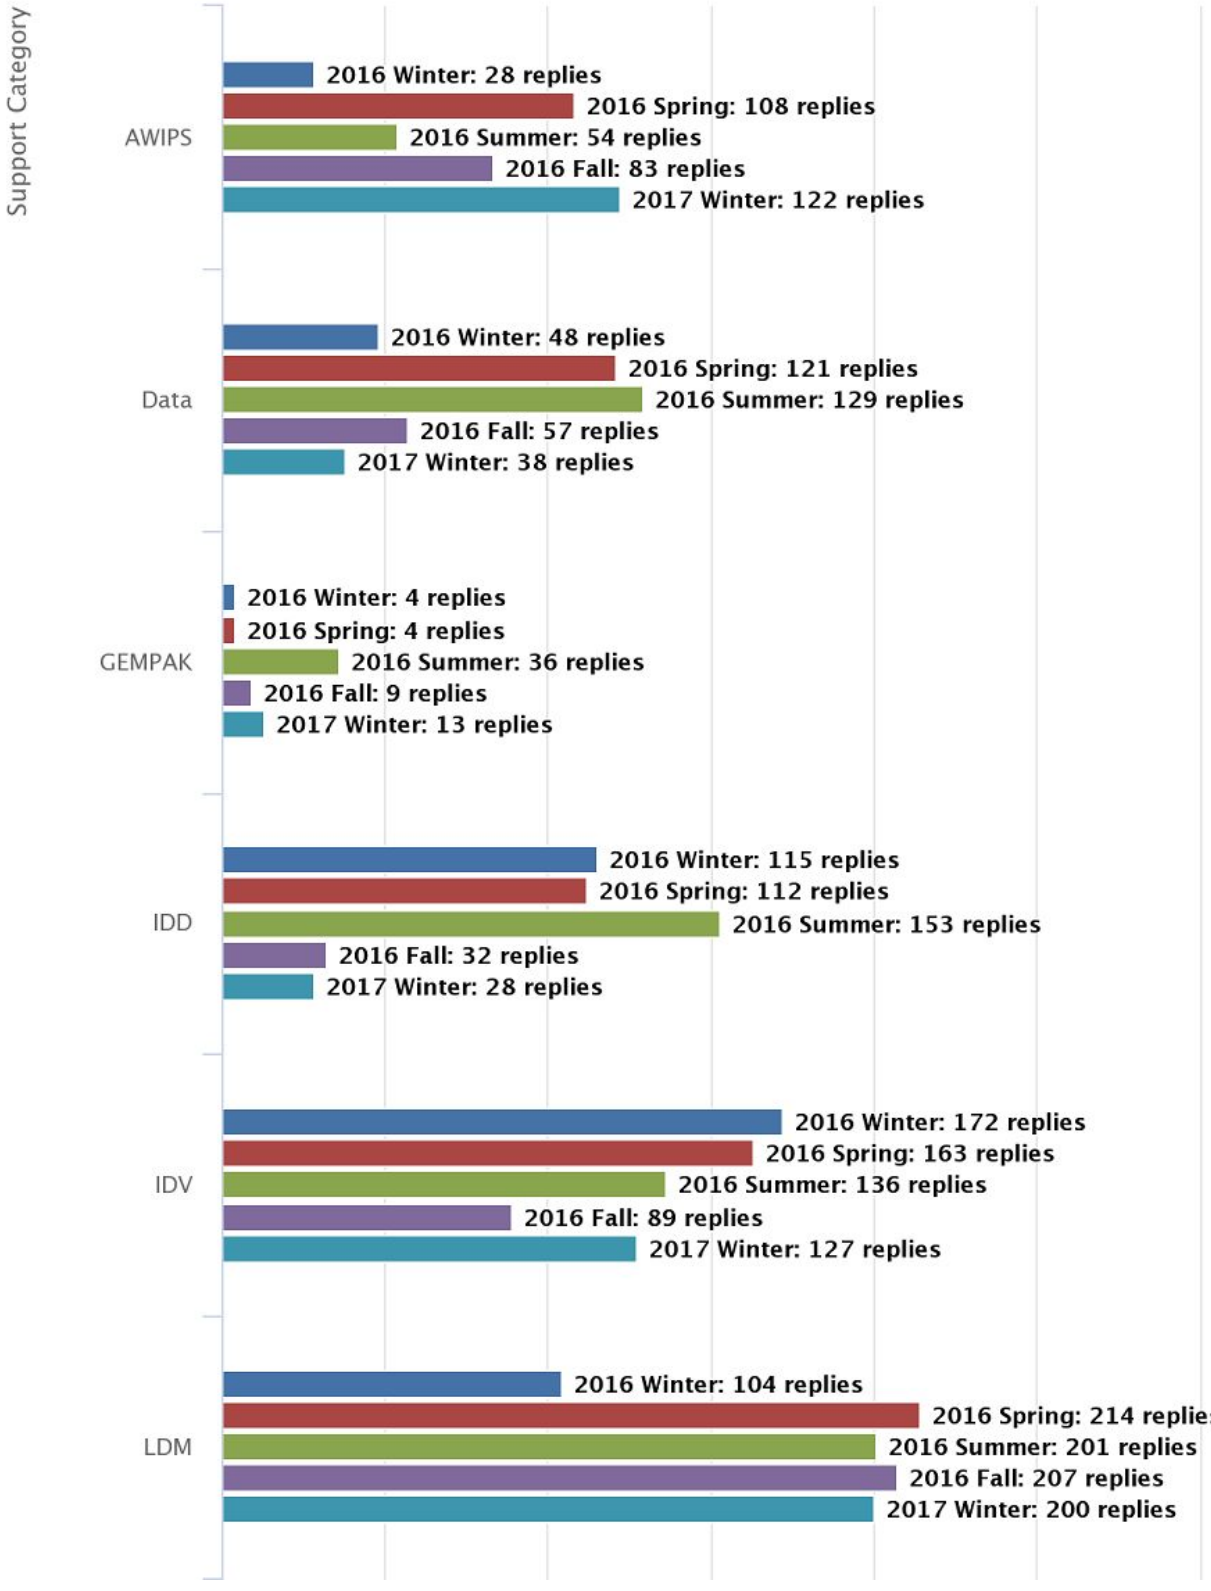

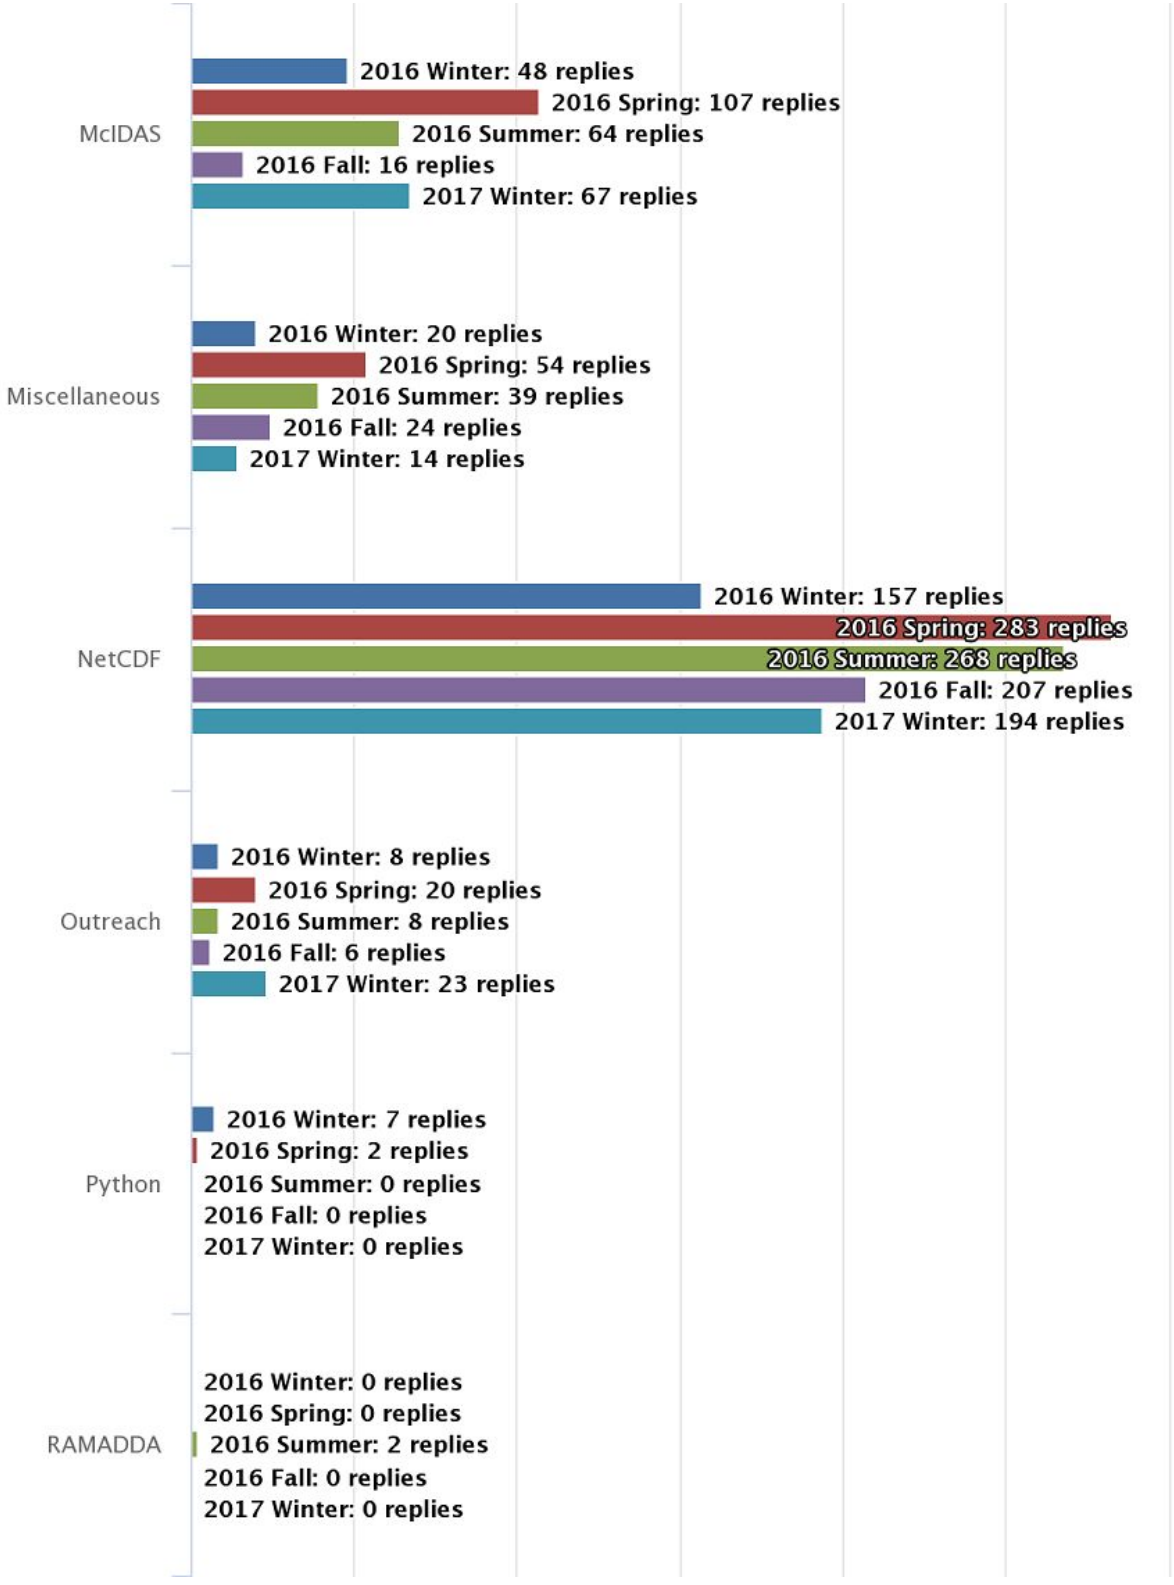
McIDAS
2016 Winter: 48 replies
2016 Spring: 107 replies
2016 Summer: 64 replies
2016 Fall: 16 replies
2017 Winter: 67 replies
Miscellaneous
2016 Winter: 20 replies
2016 Spring: 54 replies
2016 Summer: 39 replies
2016 Fall: 24 replies
2017 Winter: 14 replies
NetCDF
2016 Winter: 157 replies
2016 Spring: 283 replies
2016 Summer: 268 replies
2016 Fall: 207 replies
2017 Winter: 194 replies
Outreach
2016 Winter: 8 replies
2016 Spring: 20 replies
2016 Summer: 8 replies
2016 Fall: 6 replies
2017 Winter: 23 replies
Python
2016 Winter: 7 replies
2016 Spring: 2 replies
2016 Summer: 0 replies
2016 Fall: 0 replies
2017 Winter: 0 replies
RAMADDA
2016 Winter: 0 replies
2016 Spring: 0 replies
2016 Summer: 2 replies
2016 Fall: 0 replies
2017 Winter: 0 replies

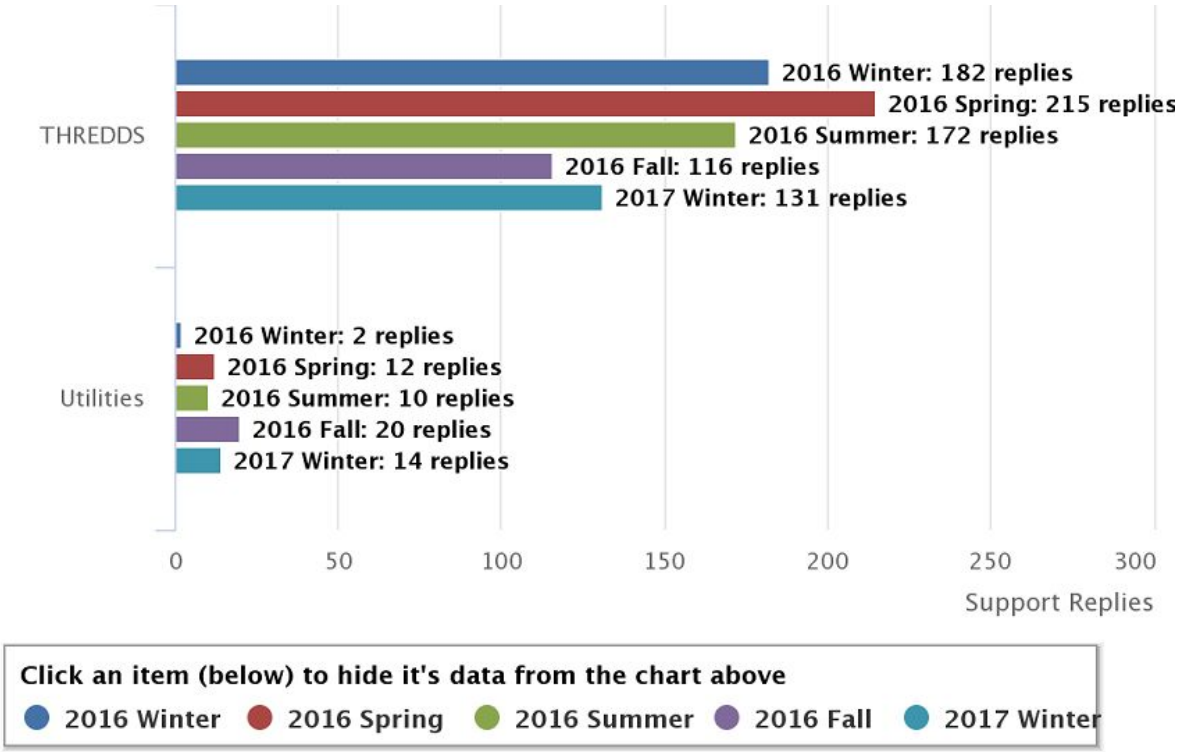

Individual support activities included in the categories shown above are listed in the following table.

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 continues to be substantial.
- Support for AWIPS has been increasing steadily and now exceeds that for GEMPAK.
- 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.

Additional User Support Metrics

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 *April 8, 2017*

Status Report: THREDDS

October 2016 - April 2017

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

Activities Since the Last Status Report

The THREDDS Project

The THREDDS Project encompases 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 4.6.8 (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

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 beginning of November. 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 end of May...*real soon now*™

Progress has been made on the following:

- 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.
- Working to integrate ncSOS 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
- John Caron is now employed by Google, and as such will have minimal to no involvement with future TDS development. Prior to his employment with Google, John was going to continue to extend TDS on a contract basis.

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
- Working to support automated transformation of output from electronic animal tagging datasets via Rosetta.

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 *potentially* extend its functionality to interface with the AWIPS-II EDEX server
- Continue to implement a Rosetta interface for each discrete sampling geometry (DSG)

from the CF-1.6 specification (http://cfconventions.org/Data/cf-conventions/cf-conventions-1.6/build/cf-conventions.html#discrete-sampling-geometries) Only the profile DSG is left to implement.

The following active proposals directly involve THREDDS work:

- The NASA ACCESS award with JPL is entering into the second half of the first 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 NASA ROSES ACCESS award: "High Performance Multidisciplinary Open Standard Data Services to Serve Terrestrial Environmental Modeling" with USGS CIDA.

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]

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

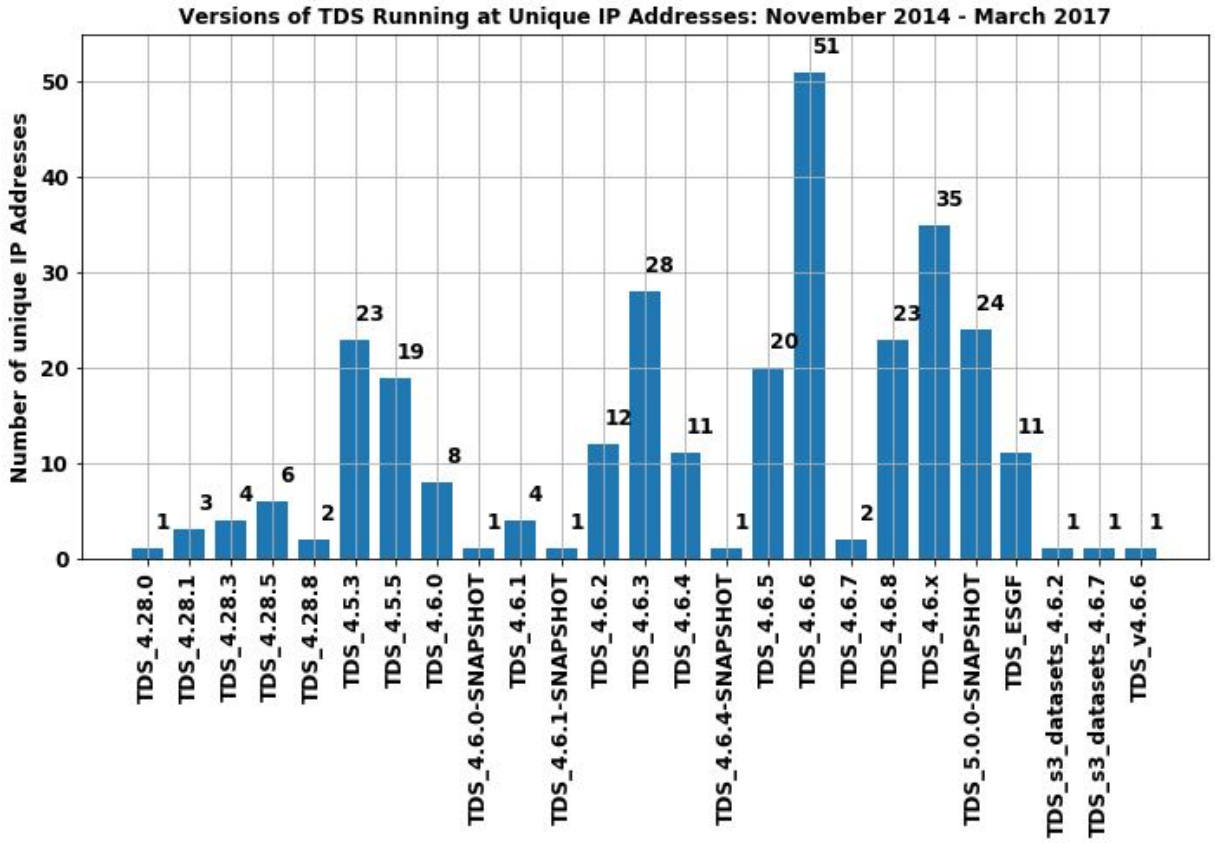

5,826 unique IPs started up thredds between November 2014 and March 2017, 293 of which are publicly accessible servers. This information is only known for servers running v4.5.3 and above. The differences in these numbers could be due to:

- 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 publically reachable
- 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 (v4.6.8 was the most current release during this period, and was released on 6 January 2017). 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, 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.

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 *March 2017*