Unidata 2013
Information Technology Survey Results

February 2013

Unidata Program Center
UCAR Office of Programs
P.O. Box 3000
Boulder, CO 80307-3000
www.unidata.ucar.edu

Mohan Ramamurthy, Director

1 Introduction

In early 2013, the Unidata Program Center surveyed members of the Unidata academic community to form a picture of the Information Technology / Computing Infrastructure situation in Unidata member university departments. This report summarizes the results of that survey.

The survey was conducted via a web form on the Unidata web site (http://www.unidata.ucar.edu/community/surveys/2013it/survey.html). An invitation to participate was sent to roughly 1100 e-mail addresses on Unidata's 'educommunity' mailing list, which consists of members of the larger 'community' mailing list who have e-mail addresses ending in '.edu'. The survey was also publicized on the Unidata web site itself. In total, we received 33 responses to the invitation

The survey itself consisted of 11 multiple-choice questions regarding information technology support and technologies in use, along with 4 questions about specific Unidata technologies and 3 open-ended questions. Respondents were allowed to leave any question blank, and to select multiple answers to the multiple-choice questions.

This document contains a brief summary of the responses to the survey, stripped of identifying data.

2 Demographic Profile

Although the survey allowed anonymous submissions, most respondents chose to provide information about themselves. In addition to names and departmental information, respondents were asked to identify themselves as one of the following:

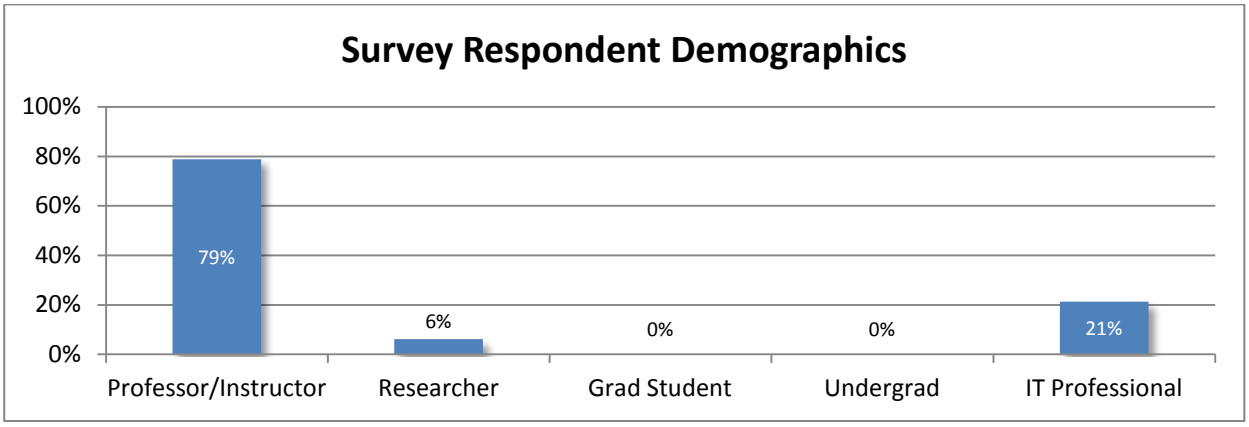

3 Department IT Staffing

A portion of the survey concerned the nature of information technology resources used by the respondent's department.

We asked: Does your department have a dedicated information technology staff, or do you use campus IT resources?

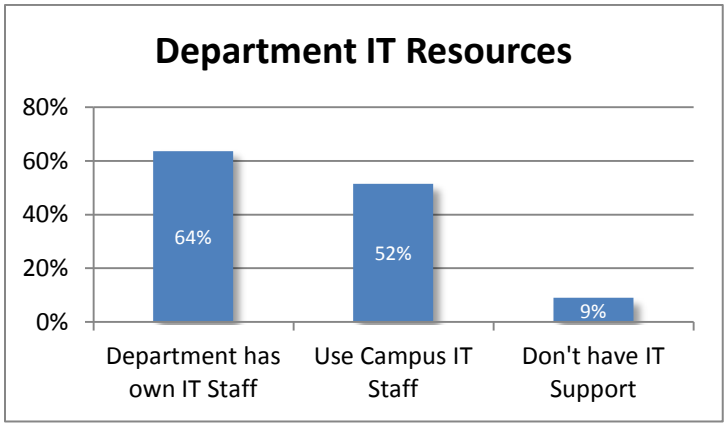

With the follow-on question: If your department has its own IT group, does the staff have a background related to your department's science focus?

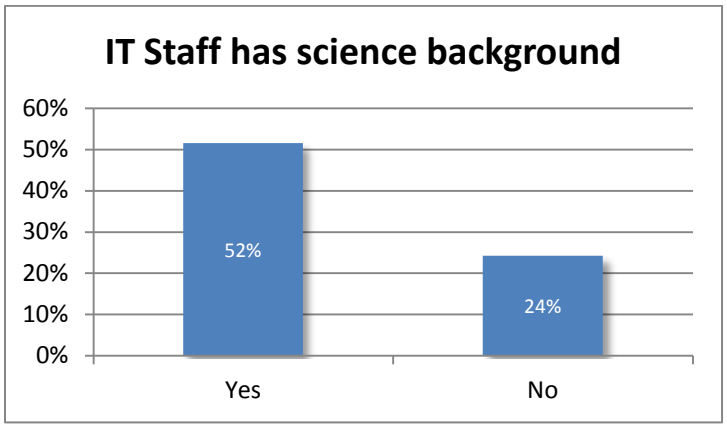

4 Computing Support

A portion of the survey concerned who provides computing support for various types of routine, scientific, and classroom activities.

We asked:
Who provides support for "desktop computing" — things like electronic mail, web browsing, and office applications — in your department?

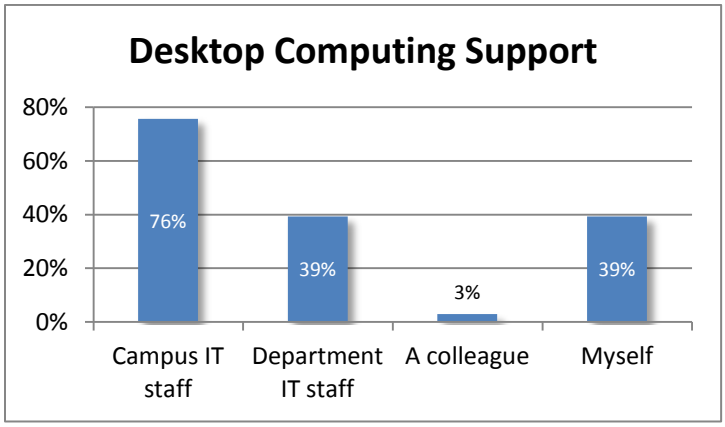

We asked:
Who provides support for "scientific computing" — things like analysis tools, visualization tools, and software development environments — in your department?

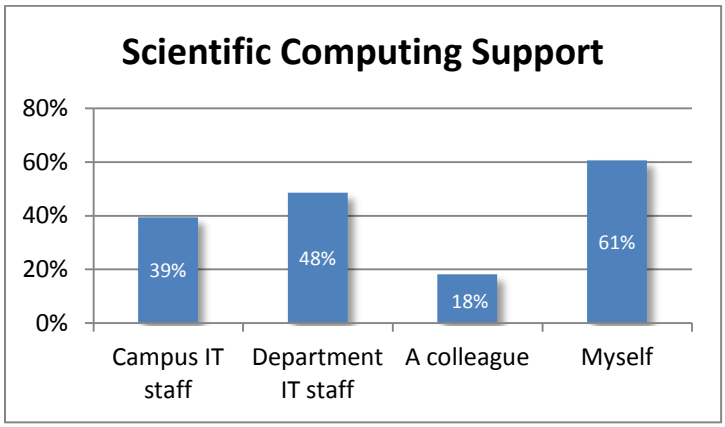

We asked:
Who provides support for "data infrastructure" — things like retrieving or accessing remote data sets, or designing ways to store data locally — in your department?

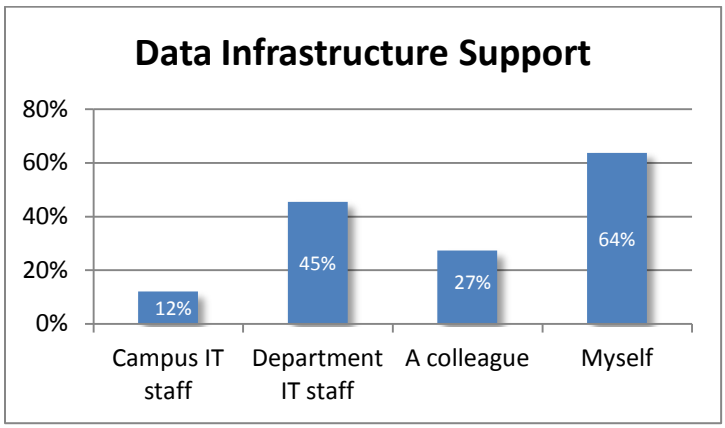

5 Student Support and Resources

A portion of the survey concered the types of computing resources students typically use in class and in their own work and who supports those resources.

We asked:
Who provides support for student computer labs or other student-facing computing resources?

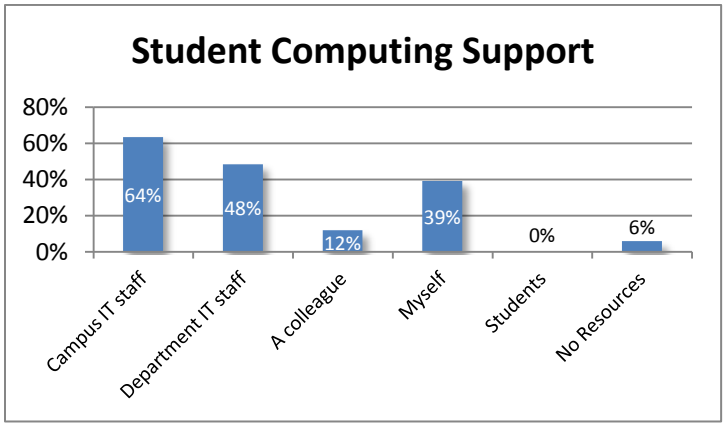

We asked:
Do students typically use computing resources in the classroom or lab environment in your department?

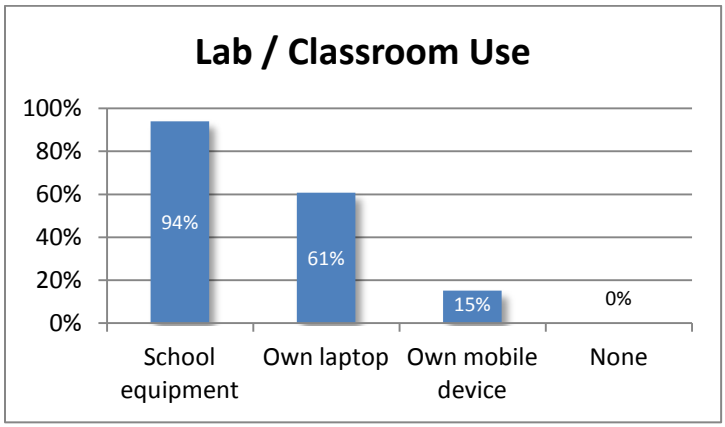

We asked:
What type of computing resources do students typically use when doing coursework for classes in your department?

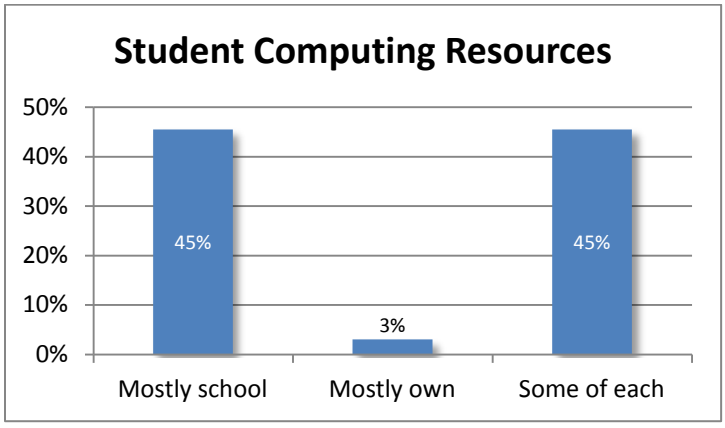

6 Platforms in Use

A portion of the survey concerned the computing platforms / operating systems in use for scientific computing and data infrastructure.

We asked:
What are the predominant platforms for scientific computing applications (analysis tools, visualization tools, and software development environments) in your department?

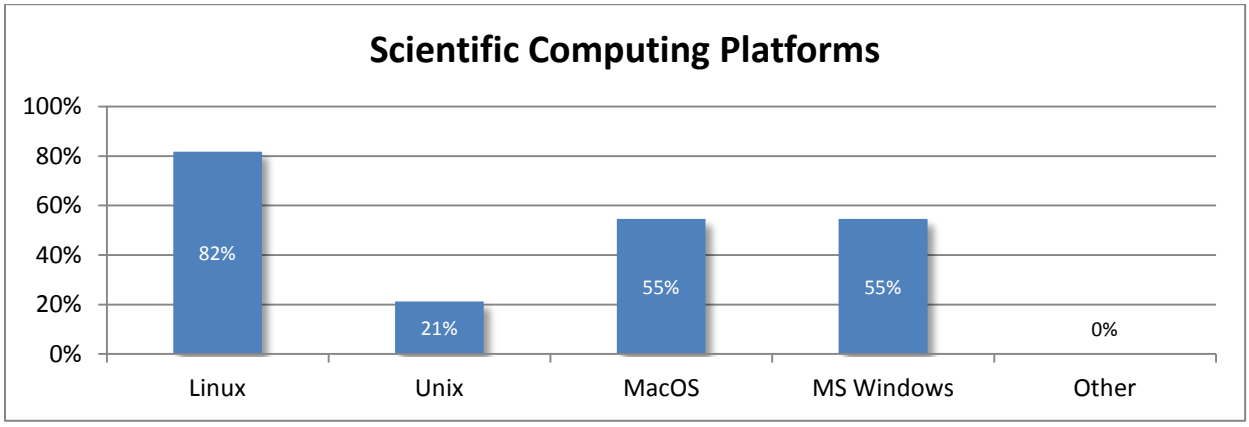

We asked:
What are the predominant platforms for data infrastructure computing applications in your department?

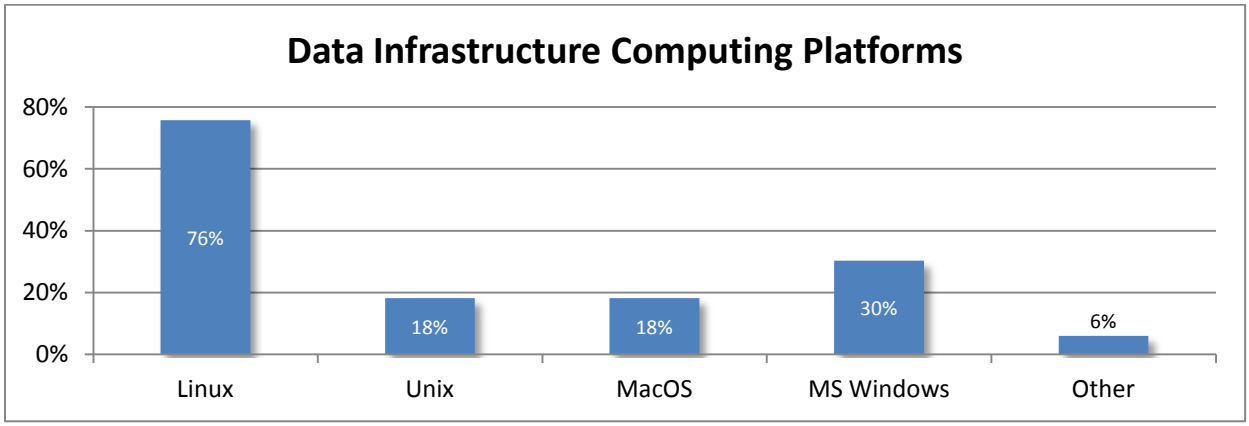

7 Technologies in Use

A portion of the survey concerned the types of scientific computing technologies in use in the department.

We asked:
Which scientific computing technologies are in use in your department?

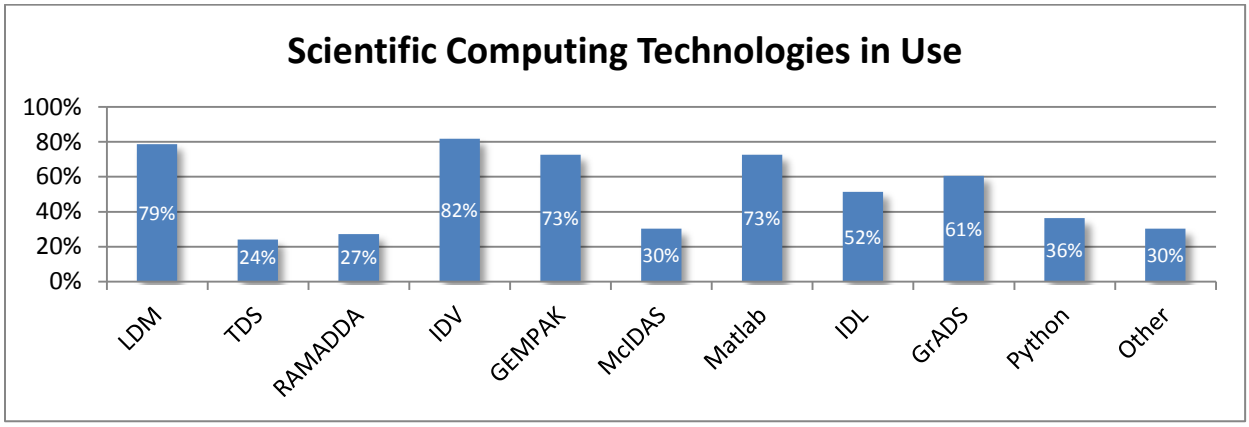

We asked:
Does your department operate any data servers, either for internal departmental use or for use by the wider community?

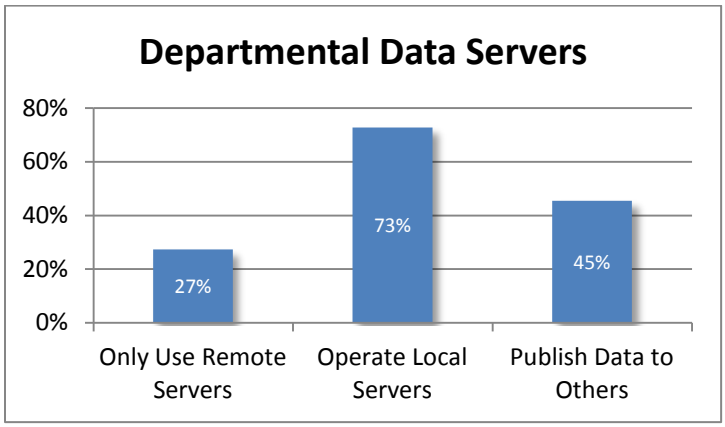

8 Use of Unidata's Local Data Manager

A portion of the survey concerned use of Unidata's Local Data Manager (LDM) software and the associated Internet Data Distribution (IDD) network. We asked:

Please tell us about your department's use of the LDM software and the IDD network. If you do not use the LDM or participate in the IDD network, can you tell us why you choose not to? Are IT policies or the availability of IT resources a factor in the decision?

We received the following responses (stripped of identifying information):

I am running LDM on one of my boxes. The data are archived and the disk is virtually mounted on the campus server where Unidata visualization software is installed.

Campus IT recently set up an LDM machine to serve as the focus or gateway for the NWC and then the individual groups within the NWC are to pull from that server. The ideas was to limit the number of machine pulling from the upstream site to a smaller number. We make use of all the various feed types available on the IDD.

We are the sole distributor of NLDN data to other universities. We ingest most LDM feeds, mostly from higher-level relays of motherlode. We also ingest feeds from non-Motherlode sources, e.g. Iowa St. and Texas Tech Mesonets, HRRR, USPLN. We feed motherlode-sourced datastreams to sites as well.

We use the LDM to fetch data from an upstream host.

Our department currently runs an LDM but it is only used by a few people. We are hoping to beef up our use of LDM in the near future.

We use the LDM primarily with our own NOAAport receiver system (Planetary Data) and the IDD for a relatively small number of products. Our network infrastructure (WAN) could not support the necessary bandwidth for the data products we need.

I believe we use LDM to receive our data from Unidata and to transport it to others. We also have used it to send our own model output to the NWS office nearby.

Weather observations and various model grids are obtained through LDM. Some observations are archived. forever, but most data cycle in and out over a period of ~1 week. Our website (weather.uwyo.edu) triggers gempak scripts to visualize the data.

This is managed by a colleague who is aware of policies.

We use it to collect and distribute data products from/to other institutions.

We are a Tier-1 IDD relay and a THREDDS and RAMADDA server site

We have an operational LDM, but are not using it as much as IDV can access datasets remotely. We have our LDM running with the default feed. We have plans to be more selective about the data we download, and to better tailor the IDV to access this data, but we never get around to it. We do have our RAMADDA set up to serve the data downloaded via the LDM, but it is not organized in a way that students can access it as easily, like motherlode is.

We use LDM extensively for data ingest for our website and on 3-4 servers total. We use it for NOAAPORT ingest to relay into the IDD. Without LDM, our program would be non-existent.

Use services for providing specific data sets to feed project related data needs.

Another faculty handles our computer lab software.

I don't know the details but do know that we're a unidata distribution node.

We manage our NOAA port data stream with our LDM

We use the LDM to both ingest data from our NOAAPort receivers and from IDD sources. We use the LDM to share data within the campus environment. We use the LDM to share images that we generate internally on different machines to then have them available on the web server. Our IT networking folks have been very supportive of our efforts and help to monitor our bandwidth usage and availability.

LDM used on Dept server for collection of real-time observations and model output(NAM, GFS, etc). Some data is stored indefinitely, especially for potential case studies(for class and research).

We run the LDM, which mostly serves our Linux lab that is used primarily by our undergraduate students.

We do not use the LDM or participate in the IDD network. Our campus IT policy is to support Windows only. Linux or Unix is not supported (some of us have protested this policy). That is why we don't use the LDM.

We have a 3-month LDM archive, in service of an old course that used GARP software. I recently updated the course to IDV, and find the highly folderized 3 month LDM archive of .gem files inscrutable, completely useless. Indeed, pointing IDV at it now, I often can't even use it, as I get this error: "Too many dataset nodes found in the catalog..." We need to update all this.

*We do use the IDD network and LDM software. Doing this as an "end leaf" is no problem here. However, doing this to be a relay to other down stream sights has problems, as the servers live behind a firewall. *Just* recently, I placed one of my data servers in a "DMZ" outside the firewall zone. It was many months of work with our IT folks and support from my Dean to arrange this. Now I have the possibility of feed down stream hosts. However, one more issue - I run on a Mac which has a known problem with LDM that has not been fixed. I work around this by running a script that kills off hung up LDM jobs, cleans up LDM, deletes and makes a new queue and starts LDM anew. This done every two hours. This is not the best environment for an upstream host on the IDD. We appreciate being able to be on the IDD as an "end-leaf" and able to get this data to our students. As a note, use of Linux is possible, but there are no available resources or computing capacity to support me at this time, nor funding to have our own Linux server system - which would be the right answer here.*

We use the LDM/IDD as a downstream node. The LDM requests and saves selected data, and triggers scripts that generate GIF images that are served to the public by a Web server.

We are a relay node for NOAAport and some Level 2 data. We use it for our own purposes, and share the data with others as we are able to.

We import on the order of 15 GB per day of several data streams for use in our meteorology computer lab, as well as some automatically updated forecast products for our departmental website.

We run a number of LDM servers for various functions, support of lab, run data servers, distribute satellite downlink data, etc.

9 Use of Unidata's THREDDS Data Server

A portion of the survey concerned use of Unidata's THREDDS Data Server (TDS) software. We asked:

Please tell us about your department's use of the THREDDS Data Server (TDS). If you run a TDS, do you make data available outside your department? If you do not run a TDS, can you tell us why you choose not to? Are IT policies or the availability of IT resources a factor in the decision?

We received the following responses (stripped of identifying information):

We do not run a TDS. We just do not have the support/resources that we need to do this.

We do not user THREDDS

I have set up and deployed a test TDS server but have no immediate plans to make it operational. We use RAMADDA instead. We are in pretty good shape in terms of computer resources to run TDS, but the number of person-hours required to support two similar data serving systems do not warrant using both. Should the "lights go out at the RAMADDA inn", we would likely switch to TDS.

We haven't been that familiar with THREDDS and its benefits, so it hasn't been used.

Serving large amounts of data would have to be coordinated with our campus IT dept. as well as the larger state college system.

We set up a THREDDS server for a data archive we have. The data is available to everyone.

We experimented with a TDS and were planning big, but lack of time and interest has kept it small - we only serve our aircraft and airborne radar data internally using our TDS

Use it to gain access to outside servers.

Yes, we make our THREDDS site available to outside entities.

We had once set up a TDS to serve case study data, but we replaced this with RAMADDA. RAMADDA was easier to set up and manage.

Hoping to implement THREDDS at some point in the future, competing priorities have put this on the back burner. Faculty buy-in to the concept may be one issue preventing role-out.

Server is part of a comprehensive data store development that includes 6 nodes, one of which is THREDDS. It is used to store netCDF files form modeling efforts.

Not familiar with this.

Not a clue. Probably the sign of a good IT staff that makes everything transparent.

We have struggled to consistently run a THREDDS server, but we have used it to organize model forecast data for internal access to other platforms

We had a TDS running years ago but the management of the data catalog was not very mature at the time (according to our systems administrator who has since left). We plan on revisiting this topic this calendar year. We will eventually make the data available to the Internet community.

Do not currently see a need for one.

No

We do not run a TDS for the same reason as the last answer. For both, we have no support. Another faculty member who knows Linux keeps the machines running but that is not his primary job. Teaching classes is his job so we can't overload him more than he is already.

We just don't have the expertise to know what this means.

We do not use THREDDS. Not sure if it would be an IT issue, as we don't have time to have this.

We don't run a TDS. I tried (once on my own, another time by attending part of a Unidata workshop on TDS) to figure out how to set up and configure one, but haven't completed the process. That might reflect the fact that I haven't yet understood fully whether we really need one or not. (We aren't a big dataset provider--really just data from the weather station on our roof.)

We run threads server for lab. Other units in school run threads server for climate data dist.

10 Use of RAMADDA

A portion of the survey concerned use of the RAMADDA Scientific Content Management system, originally developed at Unidata. (RAMADDA is now an Open Source project; the Unidata Program Center staff provides some support.) We asked:

Please tell us about your department's use of the RAMADDA scientific content managment system. If you run a RAMADDA server, do you make data available outside your department? If you do not run a RAMADDA server, can you tell us why you choose not to? Are IT policies or the availability of IT resources a factor in the decision?

We received the following responses (stripped of identifying information):

See my TDS comments above. (We do not run a TDS. We just do not have the support/resources that we need to do this.)

We use our RAMADDA server to serve data from the IDD to our users and also to make available specific teaching datasets. In addition, some reference resources are served as well. We find the ease of setup and accessibility to be an important feature in us choosing RAMADDA.

I have been actively following RAMADDA development and deployment for several years and remain in close contact with the developers as they continue work on it. We currently have a ramadda server that is available to all external users. We serve out a subset of the CFSR, recent hours of the HRRR, regional WRF model output, and several IDV bundles via this server. I am using RAMADDA in my junior-level computing application class for the students to upload/download IDV bundles as well as other datasets.

We haven't yet set up a RAMADDA server, but this is in the plans for the near future.

We have not identified any need for RAMADDA at this time.

I believe we now have a RAMADDA system also set up for our archive, and it should be available to others outside our department.

Yes, we make our RAMADDA site data available to outside entities

We are running RAMADDA and it is being used in a variety of classes - Synoptic, Instruments, NWP. The integration with IDV is a big plus. Content is accessible outside the department, but we haven't really advertised this in any way.

Hoping to implement RAMADDA at some point in the future, competing priorities have put this on the back burner.

RAMADDA is part of the same data store, used to store files of all sorts. These are normally arranged by project.

Not familiar with this.

Never heard of RAMADDA

We don't use this

We have simply been too busy to explore this.

Do not currently see a need for one.

We are increasingly using RAMADDA servers at other institutions, and would like to set up our own

We do not use RAMADDA.

I just started a RAMADDA server. I love it.

We do not use RAMADDA, and not sure if it would be an IT issue, more a time issue to work on having this.

We use RAMADDA to serve data acquired via the LDM/IDD, often to local users of the IDV (mostly students, sometimes faculty). We also make some IDV bundles available via the RAMADDA server.

I'd be interested to learn more about using RAMADDA's capabilities more creatively, though, including for classroom use.

We run ramadda for the computer lab. We also access Mother Lode

11 Use of Visualization Software

A portion of the survey concerned use of Unidata-supplied visualization software. We asked:

Please tell us about your department's use of Unidata-supplied visualization software (the IDV, GEMPAK, and McIDAS). If you use any of these packages, can you provide details? If you use non-Unidata-supplied visualization software, what packages do you use, and how did you make the choice? Are IT policies or the availability of IT resources a factor in the decision?

We received the following responses (stripped of identifying information):

We have installed IDV throughout the campus wide computer labs. GEMPAK is used in coursework here (mostly synoptic met) and is installed on the campus server.

We are still a GEMPAK shop for much of our visualization but also introduce the IDV as an alternate viewer. GEMPAK is more in teaching and the IDV seems to have found a niche in research and use as a discovery tool. Python has been gaining a huge acceptance and is being used as the primary visualization tool used by the student run Weather Lab. Python was recently introduced in our senior synoptic lab and the department is teaching for the first time a computer science programming course based on Python,

Different PI's and their grad students use different visualization packages. GEMPAK, NCL, and Matlab are most widely used. IDV use remains somewhat limited, but I am now "seeding" our undergrads by using it actively in my class ... several of the students have already downloaded it on their own computers. McIDAS-X is still used by one of our PI's. It will be curious to see how his research team reacts to Mc-V. I suspect they will use legacy Mc-X as long as they can. We have no IT policies that favor/frown on one package vs. another ... although preference is strongly given to open-source packages. I am sure it is not unique to our department, but there is a lot of inertia when it comes to what tools a particular PI and his/her team prefers to use. "If it ain't broke don't fix it" applies very well.

We use IDV and GEMPAK for weather discussions and lab exercises.

The use of these software packages varies a lot from one research group to another. The groups that use real-time data and do weather-related research do use GEMPAK

and IDV, but other groups primarily use other resources such as NCL, GrADS, IDL, and Python. IDV is used some in the synoptic and mesoscale courses. GEMPAK is primarily used by people who learned it somewhere else (it's not taught to new students, as far as I know.) Given the research focus of the department, the software packages that make high-quality, publication-quality figures are a high priority.

Gempak is still our primary tool for the classroom. McIDAS tends to be more difficult to navigate for students to start out with and we only use it on our LDM server to provide a local data source for the IDV. We use the IDV in a more limited role in the classroom. In addition, we use Grads, Solo3, GRlevelx, Bufkit, RAOB and probably some others.

We traditionally relied almost entirely on gempak (and garp) and still use these for most of our visualization. Nmap2 is used a little. We dabble with IDV but have not found it good enough to convert existing labs or scripts generating web images from gempak to IDV.

IDV has become the undergraduate choice, while graduate students are still trained to use gempak and to maintain their scripts. The reason probably is that our faculty have not delved deep enough into IDV/Jython to derive products within IDV. Several graduate students now use NCL for data analysis.

These are used in classroom and research projects.

Primary developer of McIDAS package.

I use GEMPAK all the time, and IDV occasionally (it's still far too slow to be practically useful to me.) GEMPAK still has the best and most efficient way to quickly visualize data of any packages. I also use NCL for data manipulation and plotting WRF grids. I know GEMPAK development has stagnated, but GEMPAK reading WRF grids and plotting them (or at least an official WRF2GEM utility) would be great.

We routinely make use of ALL of these packages. We plan to use the AWIPS-2 packages as well as they mature and become available.

We are using IDV as the primary tool for case study analysis in Synoptic. Students also use it to create case study presentations in Mesoscale. Students use GEMPAK for research. We do not use McIDAS - the configuration is more difficult than the other software packages.

We use GEMPAK and McIDAS for data analysis and automatic imagery generation for our website. We use IDV for a quick look at data. IDV has been buggy as we move towards more dual-screen environments. The release cycle of IDV does not line up nicely with our biannual reimaging of our 30+ lab machines, which usually puts us at least one version behind for the semester. If new versions of IDV could be released in in late July/early August instead of September, it would be a big help to our program.

only used infrequently, if then IDV

Our intro classes use Internet products.

I use Matlab. GEMPAK is so limited that I now refuse to join the graduate committee of any student who uses it because they're never able to implement one's suggestions.

Python's nice too. Not as many tools as a full Matlab install, but it's free. The install is a pain though, layers on layers on layers in order to get something like Matlab's level of scientific computing. And then the layers turn out to be semi-redundant and contradictory. What a mess.

We used GEMPAK programs extensively in many classes - it's quick and easy. We also run IDV, but its use is limited in my case to research and assignments - it's not a great operational tool. We have nearly phased out our use of McIDAS.

We use the GEMPAK family of programs (especially NMAP2 and GARP) in a number of our courses for both meteorology majors and for some of our courses suited for the minor. Some of us use the IDV more than others, while there is discussion of migrating from GEMPAK to IDV due to uncertainty in GEMPAK applications in the future.

Gempak is used for weather forecasting by myself (I'm the forecaster for the campus). It's also a tool for teaching a synoptic meteorology course. I administer the software (updating for example). I have installed IDV on our system, but I haven't gotten around to getting familiar with it yet (planned uses in the future).

We and our students use GARP for real-time forecasting and research purposes. We use IDV for more research purposes. We use some other software like GR2Analyst for radar data visualization.

We use GEMPAK/N-AWIPS in a lab, along with IDV. This software is also loaded on many individual research workstations for grad student and faculty use.

I use IDV in the classroom and for my own researh

I use GRADS extensively

We do use the IDV for visualization in several classes. Mostly the upper division meteorology students use the IDV for class assignments and occasionally student research. We introduce the IDV in the introductory meteorology classes so they will be familiar with the package.

I use IDV, and there may be a research group that uses GEMPAK or McIDAS.

We do use McIDAS-X and IDV/McIDAS-V. McIDAS-X is used a bit more as I am more familiar with it, than IDV/McIDAS-V. McIDAS-X drive a routine looper style display in the main classroom that Weather and Climate is taught in. Both McIDAS-X and IDV/McIDAS-V are installed on all department laptops, which are used in my Weather and Climate classes roughly 3 or so times per semester...and growing. IDV/McIDAS-V will be used in an outreach project in the coming semester.

We started in the mid-1990s using WXP, with support from Unidata initially and then with the support of the developer and the community of WXP users (after Unidata decided not to support WXP directly). Graphical products that are served by our Web server and that are generated from data acquired using LDM/IDD, are still generated automatically by scripts that run WXP. We also use scripts written locally to run WXP in our classes. However, we also use, and teach the use of, the IDV in our classes. At some point I'd like to take the time to figure out how to create IDV scripts that generate many, if not all, of the products that we currently generate for our Web server using WXP.

Students are given exposure to these packages in several of our meteorology courses, including one that is specifically designed for such a purpose.

We use IDV for classes. Gempak and McIDAS are used to generate graphics for our web servers.

12 Challenges

The final question of the survey asked:

What are the biggest challenges facing your department with regard to scientific computing, data acquisition, and data management? How does the availability of IT resources affect these challenges? How can Unidata help?

We received the following responses (stripped of identifying information):

The big struggle is determining whether I want to buy my own boxes or contribute support to the campus IT folks in support of their soft/hard infrastructure. Given that we do not have IT support within my dept. the campus IT is a viable option (less hassle!).

As systems and software become more complex, knowledge and support duties of IT staff is more than a one person job. Especially when including the additional duties that fall to the IT support of the department in the way of help desks, maintenance of audit and IT policies and other software support. In addition, students at all levels are seemingly less competent with respect to analytic computing skills (programming, debugging, compiling third party code, etc). This is a challenge in its own right.

We are of course very fortunate here to have two state-funded people to handle in-house IT needs for the department. The fact that we are both meteorologists by training certainly helps, and I don't know how pervasive this is at other universities. Yet as our department grows, we feel it is getting more difficult to spread our resources as we once did. We are now partnering with our research IT group on campus to house a computing cluster for four of our new faculty. The folks who manage the research IT cluster do not have experience with meteorological software, and understandably prefer to install pre-built binaries (e.g. via rpm). The recent release of GEMPAK in rpm format is greatly appreciated, for that reason!
A major challenge going forward is educating our faculty and grad students about the tools that Unidata (as well as other providers) give us. It really takes serious effort (and time) from busy IT folks like us, to inform our equally busy core users about what servers or packages like RAMADDA, TDS or IDV provide. The increasing tendency, as alluded to above, is for more and more inertia, where folks rely on the tools they have used in the past, simply because they know they can get them to work, even if they are not the most efficient way to do things. Unidata needs to give some thought to greasing the path, so as to remove this inertia.

We have a part-time IT person who is assigned to us for 25% of his work time.

I would say the biggest challenges are:
1) Collecting, managing, and analyzing very large datasets (such as high-resolution model output or long-term reanalyses)

2) Getting students to acquire the necessary skills to work with these datasets effectively and efficiently.

Keeping up with the storage needs and network bandwidth is always a concern. Linux is a great platform but also is a moving target as far as getting things to work properly - some applications want the latest version of the system libraries and utilities while others seem to be developed on older releases. Getting anything to compile and link properly with netcdf is a frequent source of frustration.

No one has enough time to learn new things. We cannot keep up with the model output we generate, and we rapidly fill our servers. I suppose a Unidata-administered cloud where we could store the terrabytes of data would be wonderful but that is probably not possible. We always encounter inertia with having to convert data formats, or in spinning up new students with visualization software capable of working with our data formats.

We are fortunate to have good IT support with an interest in atmospheric science. Unidata remains very valuable.

Money, time, and IT support

Using Web Services (Map, Feature and Coverage) to distribute products to users.

Biggest challenge here is teaching people to use the tools that are available to them. We have sufficient IT resources, but insufficient experience with the tools. This situation is improving with time, but still has a long way to go. I'd love to get people (including myself) using IDV more often, but it's too frustrating to regularly use in its current form.

BACKUPS and Archival storage as required by NFS and other funding agencies! Penn State is upgrading its backbone for research/data to 100 GB and 10 GB throughput, so that should be enough bandwith for a while.

We have IT support at the college level, which is one up from the department. It is only a few people and they are stretched very thin. Only one person has familiarity with some Unidata products - mainly installation. If we want IT to do it, new versions are only installed once per year in the summer. Our main challenge is that all software and data configuration and customization must be done by the faculty.

Funding for keeping classroom/lab PCs updated is limited as well as server is limited. Time is also a factor, I attended the THREDDS training course 2+ years ago, still haven't found time to implement.

bandwidth, infrastrucutre such as switches, upgrading to higherthroughput setup, the "last mile" is antiquated, problematic: working with University admin and IT department to get system up to date.
Firewalls, security settings around campus, IP address assignments, server settings and configurations for use in campus networks,
Not sure UNIDATA can help on these ...

Training, time to update our skills.

Data movement and archive architecture are a failure at this point. We build archives as snapshots one period (often a day or hour) at a time because that's how the data comes in. If you do synoptic (who does nowadays?) that's wonderful. If you do time series analysis (climate, climate change, statistical forecasting, all the standard stuff these days), you're dead in the water. Why, because you have to download every last file in the archive just to access the time series from one point. But no, you're only allowed to download a few files at a time. No such thing as a database query in meteorology. Nope, we don't store data in databases. We just cram flat files in folders within folders and call it organized. 90% of my effort goes into getting the data, 9% into reading it and 1% into doing science. This is just plain wrong.

Simply keeping everything running well even with our IT support has been a struggle. Unidata could help by continuing its level of support and simplifying the manner in which students access/manipulate information.
I would like to see more shared resources within the community and integration of our software into other platforms such as ArcGIS.

We find that we have trouble with data management (aging of old products in a fast and systematic way). We often have data coming into our systems that we don't know we are getting. Sometimes we have issues with our still using older releases of the LDM, decoders and display softwares. As we move more towards the use of SQL enabled databases we find we have an even greater gap in existing skills versus projected needs.
Our hardware needs seem to grow faster than our ability to get internal funds allow for. The need for enterprise class hardware to support the data processing and data serving needs has elevated the costs of computing significantly.

Accessing remote servers with various data sets (for teaching) may be a better choice than locally collecting the data (near real-time). This is especially true of oceanographic data sets. However, setting that up is currently beyond my understanding (limited Dept IT support for this).

The ability to store, analyze, transfer, manage, and visualize large datasets, such as high resolution NWP model output.

The biggest challenge has been lack of a dedicated support specialist (mostly everything is done at the University level with a lower priority) and department politics. That environment is now changing with a chair who has left and a new exciting merger imminent with the Dept. of Geology. It has already been indicated to me that more resources will be coming with this merger and meteorology is something the geologists are excited about.

Our biggest challenge is simply to keep the computer lab running. We have 10 Linux-based PCs that run off a local server. Our faculty member who services them is hard-pressed to keep up with machines that stop working or need updates (or both). Support through a local technician would be the best answer and that is something that our College must provide. Unidata can't help us with personnel issues.

In the end, reliable management of these systems falls more and more on scientists, professors. We just don't have the kind of professional support with scientific fluency that perhaps used to be the norm.
That means Unidata needs to explain its computerese to scientists, not science to

software professionals. I find many frustrations along the road to "power user" status, and many ad hoc fixes of one inscrutable problem at a time that often don't add up to enriching my knowledge of computers and software (which are not uninteresting realms, actually).

Two challenges I run into: Time to work on things (given that there is limited support), and not sure how I'm going to upgrade/life cycle replace the server computer system that I got on our Unidata equipment grant. It will be challenging in the current budget climate to do that (although we are working on ways to arrange that), and not sure I would qualify for another Unidata equipment grant. Having more time to work on things would require me to move from Adjunct Faculty status to full time status, which is not likely also due to the budget climate.

Since the mid-90s I've (off and on) spent a great deal of time learning to use WXP and the IDV (and to a lesser extent the LDM and RAMADDA), and some aspects of these packages are still challenging for me to understand and adapt to our local needs (which tend to be mostly instructional and only occasionally research). Although I intermittently forget how to do some things that I don't do regularly and have to relearn them, I usually figure out what I need to do if I can make the time for it. My colleagues, on the other hand, haven't invested anywhere remotely as much time in mounting those learning curves, so when they do try to use something like the IDV, they don't typically persist for very long before giving up. Hence, I've had trouble nudging them toward regular use of the IDV in their classes--they instead grab what they need from the Web or use simpler, single-purpose (often OS dependent) software. I don't think they have a clue about RAMADDA (yet), even though we run a RAMADDA server. (Interestingly, some of what they grab from the Web is created by people--including us--using Unidata software.)

Money/resources. Plain and simple.

Right now, things are running fairly smoothly, and we are satisfied with the level of support we have received from Unidata. More MATLAB support and automated to/from netCDF conversion packages would be nice, or an online tutorial on this. Also, would more support for GIS-related applications be possible?

Difficult question of course. Unidata can continue to make applications more user friendly, improve user interface, improve data formatting and compatibility tools, generally reduce the friction associated with data processing in doing science.