

June 2012

Data Analysis Services Group - June 2012

News and Accomplishments

VAPOR Project

Project information is available at: <http://www.vapor.ucar.edu>

XD Vis Award:

Efforts on the integration of PIO with VDC continued, and we are now holding semi-weekly meetings to keep the project moving forward:

- All API changes requested by PIO developers have been completed
- The canonical PIO test driver has been updated to include support for VDC output.
- A 1D slice decomposition, similar to that used by the GHOST code, was added to the PIO test driver.
- Version 1.5.0 of PIO was released, and included support for VDC output.

The only outstanding items now are resolving a substantial memory leak in PIO, finishing preparation for benchmarking on Yellowstone, completing the integration of PIOVDC into the GHOST turbulence code.

KISTI Award:

Work began on modifying the VDC data translation library to support missing data values.

Development:

- Alan made several changes to vapor to enable interpolation of viewpoints, and to separate frame number for data time steps. Ashish used this version of vapor to develop keyframing code, providing smooth interpolation of keyframes. Based on John's observation that speed control is an important requirement, we developed algorithms to enable users to specify camera speed. Alan modified the vapor GUI to include UI for animation control, and Ashish developed the required code so we are ready to start integrating keyframing into the vapor GUI.
- Alan determined requirements for handling missing values in Python scripts. We found that there is no C-api for the Numpy masked array, so we needed to address this need by inserting Python statements into users' code. This was implemented in vapor's Python Pipeline, so that when a python script is applied to data with missing values, the data arrays are processed as Python Masked Arrays. We are looking into the problem of implementing derivative operators that operate on data with missing values..
- We learned from Rick Brownrigg that problems we are having with getWMSImage are a result of NASA Web server changes, and that we may need to work out an alternative means of obtaining images. We decided to include the capability of extracting sub-images from the vapor installed images, as a work-around for this problem.
- We continue to respond to queries on the VAPOR mailing list, and to post and fix bugs.

Administrative:

- John and Alan reviewed Song Lak Kang's (TTU) NSF CAREER proposal. Song Lak asked Alan and John to be unfunded collaborators on the proposal, and to provide VAPOR training to high school sciences teachers at an annual summer workshop that would be hosted by TTU if an award is made. Travel costs would be covered by TTU.
- Our student intern, Wes Jeanette, accepted a full time job with Seagate. We wish Wes the best of luck.
- John submitted the XD Vis quarterly report.

Education and Outreach:

- Alan wrote up new tutorial for WRF 2012 workshop, gave a dry run. The final tutorial was presented at the 13th WRF Workshop, with help from Yannick Ashish and John. The tutorial was well-received, attended by ~25 WRF users.
- A proposal to organize a topical session on visualization, submitted to the Journal of Computational Physics sponsored event: Conference on Frontiers in Computational Physics: Modeling the Earth System, was accepted. The conference will be held at NCAR in December.

Software Research Projects

The VAPOR team led and submitted an NSF BIGDATA (NSF 2012-499) mid-scale proposal with partners from C.U., U.C. Davis and COLA.


A third, and hopefully final, round of revisions was made to the feature tracking manuscript submitted to TVCG by Clyne, Mininni and Norton.

Production Visualization Services & Consulting

John produced a number of visualizations from a wind turbine simulation run by Branko Kosovic (RAL). Branko will be showing the animations at the WRF workshop in late June.

ASD Support

- Alan is providing visualization support for the ASD project of Gabi Pfister from ACD. She is looking at the predicted effects of climate change on air quality in the 21st century. Alan processed Gabi Pfister's data to create several variables that she indicated would be needed to create visualizations of various pollutants at ground level. These variables identify locations where pollutant average concentrations exceed thresholds over specified time intervals.
- Alan is also providing visualization support for Baylor Fox-Kemper and Peter Hamlington's ASD project that examines the turbulence in ocean water induced by contact with the atmosphere. Alan calculated a number of important derived variables and also obtained and converted 10 time steps of Peter Hamlington's data. However the temporal resolution of the data is too coarse to be animated, so we are investigating means of obtaining more frequent high-resolution samples over small regions.
- Yannick started investigating the particle display capabilities of both ParaView and VisIt.
- John developed data translators and produced preliminary visualizations for two ASD awardees - David Richter (NCAR), and Peter Ireland (Cornell) - using sample data sets. Both data sets were generated from turbulence codes including particle simulations. VAPOR has no built-in particle display capability. However, with the help of a translator written by Yannick, the particle data were converted in to geometry scene files, displayable by the scene display facility written by former intern, Kenny Gruchalla. A sample visualization is below:



Systems Projects

Security & Administration Projects

- Modified the KROLE certificate management daemon to renew all certs when it is started, in addition to the normal daily renewal. This made in response to the fall out of the building evacuation power down lasting almost 2 days.

File System Test Bed

- Lustre client RPMS build, Recovery of Glade after Nexus switch reboot. Built lustre client rpms (2.2) for blizzard, mirage³⁴, and nomad. Functionality test was done on nomad (10GigE)
- Installed Robinhood 2.3.3 and tested the ChangeLog mechanism to update the database w/o scan of FS Fixed fid2path error on Lustre which was caused from 1.8.6 to 2.2 upgrade against the recommendation, by rebuilding the targets in native Lustre2.2 ldiskfs format.

NWSC Planning & Installation

- Sessions with Jan and Brian on picnic test setup
- Performed picnicmgt1 troubleshoot after power cycle with Craig and Joey
- LUN build on picnic test system and power down tests. Two LUNs on picnic test setup was initialized on Friday, finished formatting and ready for GPFS. "Immediate availability Initialization(IAF) started for tdc1-1 and tdc1-2" around 15:42 on Friday. Both finished around June 16 4:22AM. (~13 hrs to initialize when two of them run together)
- Continued power-down tests on the tc00 for DCS3700 network interface behavior.
- Worked on getting ganglia, nagios, and all dependencies running on the CFDS test system.
- Created nagios scripts to check the health and drives status of the DCS3700 storage system.
- Started GLADE test rack equipment testing, read many vendor documents. Started testing if remote management of systems in troubleshooting situations would work.
- Travelled to NWSC to observe part of the installation of the production equipment. Participated in meetings with IBM personnel. Installed SSH and UCAR authentication related software on the GLADE management nodes to support fully exposed host network interfaces. Reviewed production GLADE servers for mistakes during configuration.
- Met with the Mellanox regional personnel.
- Continued tracking HUGAP project progress. It is not yet usable for production authentication and authorization purposes on GLADE resources.

System Support

Data Analysis & Visualization Clusters

- Installed the latest version of the ANSYS/FLUENT software.

GLADE Storage Cluster

- Kevin Raeder GPFS hang issue and mmfsck online. Kevin Raeder reported GPFS hang on a particular directory, which was later lifted when a job on two bluefire compute nodes were terminated by rebooting the nodes. There were no further issues for other users and Kevin later re-ran the identical job without problem in the same directory. Online mmfsck session did indicate an error but due to ~30M inodes involved on gpfs_dasg_proj3, we postponed the off-line fsck.
- Rebuilt 120G on DDN9900 which caused the SCSI error on oasis7 (~1:02am), took ~90 minutes
- Had Discussion with Si on "unlimited" dimension and compared performance issues with netCDF files. We observed that excessive repetition of "read" happens only when one the the variable is declared as "unlimited". Since identical file was used for the tests, we could potentially flip the "unlimited" declaration and do ncks processes with re-tagged one to accelerate the post processing.
- Had Discussion with Si and Joey on "read" vs "fread" difference.
- Performed troubleshoot on /glade/proj3 after Nexus down in the morning with Joey and Craig. Rebuilding quota table enabled it to be mounted later. (proj3 was the only FS refusing to mount)
- Initiated disk rebuild on 118D after SCSI error
- Performed NCKS test with netCDF-4.2.1-rc1 which was announced to have in-memory operation. Completion time on glade is still 9 times more than that of /tmp. We will continue monitoring the announcements on netCDF updates.
- Ran mmquotacheck on /glade/proj3 to fix the mangled fileset quotas after the emergency powerdown.

Data Transfer Cluster

- Exchanged e-mails and had in-person discussions with Sidd for future htar requests.
- Discussions with Si on pgf90 setup on mirage nodes. It turned out symbolic link for environment in /fs/local/bin was masking the original files in /fs/local/apps/pgi/linux86-64/11.5/bin. Resolution was made by renaming two offending files in /fs/local/bin.
- Worked with Si a bit to try and figure out possible causes of slow GridFTP transfer performance from NCAR to other sites when Globus Online was managing the transfer.

Other

- Remotely powered down all systems after ML evacuation due to Flagstaff fire.
- Monitored ML status announcements during Flagstaff fire
- Restarted DASG nodes, focusing on DDN9900 setup and associated servers oasis⁰⁻⁷. Relaunched Lustre, swift, and test GPFS on stratus servers after production nodes were ready. Remounted the LV volume for cmip5 after "lvchange -a y /dev/cmip5/lvol0" (with pvscan+lvscan)
- Wrote emergency powerdown procedures for GLADE for any future building evacuations