

May 2012

Data Analysis Services Group - May 2012

News and Accomplishments

VAPOR Project

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

TG GIG PY6 Award:

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

- A long-standing, and exceptionally large memory leak was finally tracked down in PIO
- Integration of VDC-specific mesh decomposition was integrated into the canonical PIO test driver
- A number of minor API changes requested by the PIO team were completed.

XD Vis Award:

The period of performance for the XD Vis Award has been extended until July, 2013.

KISTI Award:

Year two of the KISTI project has officially commenced. The KISTI contract was signed by both parties, and NCAR has received funds from KISTI. Work in year two will continue efforts to support to VAPOR for global and regional ocean simulation data sets.

Development:

A project plan for the 2.2 release of VAPOR was completed. The timing of the release is set to coincide with the completion of year two of the KISTI award in December.

A number of minor bugs have been uncovered and fixed as we continue efforts to migrate to the new internal data model based on the RegularGrid class object.

Alan made changes to allow Python pipeline to operate with RegularGrid data model.

Alan made a number of changes to the Python library, so that the WRF routines now operate on the original data rather than on the vertically re-interpolated data. All the differential operators were rewritten to operate on model (non-Cartesian) data.

We heard from Minsu Joh that she is having problems obtaining terrain images from the NCAR WMS server. Rick Brownrigg provided a workaround, but there may still be bug in getWMSImage that fails to return a tiff for some extents.

Specifications were developed to overhaul the ncdf2vdf conversion utility in the hope to make it substantially easier to use.

Administrative:

Annual reviews were completed.

The VAPOR team completed work on their two-year strategic plan. The intent of the plan, available from <https://wiki.ucar.edu/display/dasg/Planning>, is to look beyond incremental feature additions to the VAPOR suite of applications.

Education and Outreach:

Alan met with Cindy Bruyere, helped her create a vapor image of a hurricane she is working on. She also discussed the MMM needs for the upcoming July tutorial. She would like us to provide a basic page on their Online tutorial explaining how to create a visualization in VAPOR. She also wants a 10-slide PPT providing a walk-thru, or road map of the GUI, showing users how to navigate there.

Ashish Dhital started his Siparcs internship. He is implementing spline interpolation of viewpoints to prototype keyframe animation in vaporgui.

Alan made temporary changes to vapor's view panel to facilitate the development of keyframe animation, which will be used by Ashish Dhital during his SiParcs internship to interpolate viewpoints between keyframes.

A proposal to organize a topical session on visualization was submitted to the Journal of Computational Physics sponsored event: Conference on Frontiers in Computational Physics: Modeling the Earth System. The conference will be held at NCAR in December.

Final revisions were made to two invited book chapter manuscripts, one on VAPOR, and one on multiresolution methods. The book, High Performance Visualization, is scheduled to be published later this year.

Documentation

A user forum was created by Wes to facilitate communication among the VAPOR community.

Wes also took the opportunity to complete the migration of the VAPOR web site to drupal. We are now fully "drupalized"

Software Research Projects


The VAPOR team is responding to the NSF BIGDATA (NSF 2012-499) solicitation, leading a collaborative, mid-scale proposal with partners from U.C. Davis and COLA. The proposal, which is nearing completion, is due on June 13.

Feature Tracking:

Climate data compression:

Production Visualization Services & Consulting

Alan met with Mel Shapiro, experimented with particle traces in the ERICA simulation. Performed forward and backwards integration from positions in the center of the storm. This resulted in images such as the following:

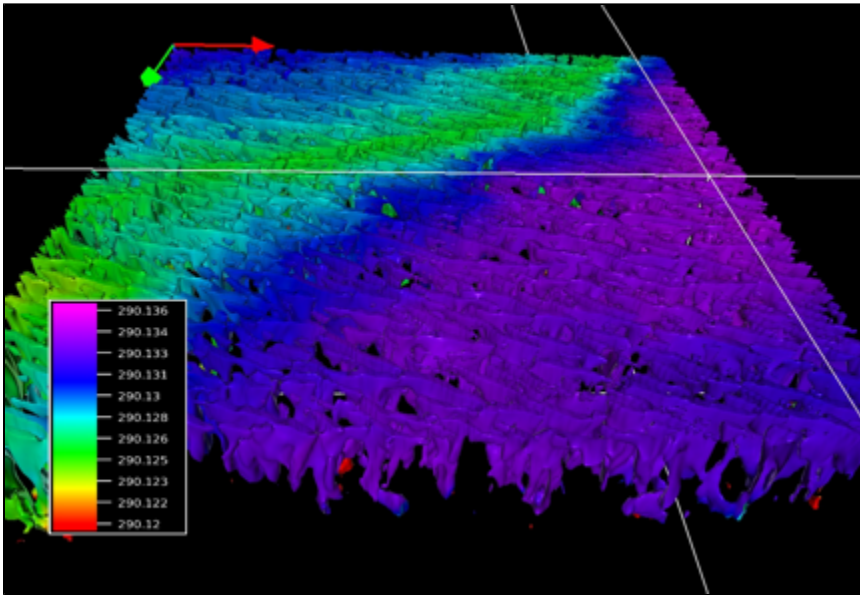


John produced a first round of visualizations for NCAR's Branko Kosovic on wind turbine array simulation data.

ASD Support

DASG staff are working with six ASD award recipients to help them visualize the data that they will be producing on Yellowstone later this summer. A key driver of this effort is the grand opening of NWSC to be held in mid October. DASG staff have been in contact with all awardees to plan the effort, and make arrangements, where possible, to gain access to preliminary example data sets.

- Alan met with Gabrielle Pfister, an NCAR awardee for the ASD project. She is heading a project that will be using WRF-CHEM. Gabrielle provided us with a month-long sample data set that should illustrate the predicted pollution at ground level (ozone, CO, and other chemicals) expected to results from climate warming, highlighting where and when the pollution exceeds allowable levels.
- Alan met with Ned Patton and Peter Sullivan regarding a Yellowstone simulation they are planning. They expressed a desire to have our assistance in developing a vapor reader for Tecplot, which produces visuals that they like a lot. We also discussed ways of visualizing their experiment in VAPOR.
- Alan obtained a data set from Peter Hamlington, on the Fox-Kemper ASD project, and visualized it in vapor. Peter is interested in turbulence patterns at the ocean surface and is working with Alan to illustrate this in VAPOR. Peter is pleased to find that VAPOR can quickly visualized phenomena that he has not previously seed. We produced images like the following isosurface of vertical velocity, colored by temperature, at the ocean surface:



NWSC Planning & Installation

- In case we get a chance to survey and identify the defective disks for NWSC system, built the xdd package and downloaded the related test scripts.
- Continued testing the xCAT software and recording observations in some wiki documentation. Tested producing localized RPM packages for installation with xCAT.
- Attended CFDS related training sessions.
- Continued work on the ingest of the HPC system authentication information from the HUGAP project. Attended meeting on distributed RAIDUS server issues.

System Support

Data Analysis & Visualization Clusters

- Following the HW manual for xw9300, tested the Power Supply, Memory DIMMs and Graphics card to find which component prevents the power-on tests on **storm3** from completing properly. Diagnosis suggests that the Motherboard should be replaced.
- Re-enabled samba on **blizzard** for visualization work
- Upgraded IDL from version 7.1 to 8.2.
- Gave privileged access for file management commands to Gary Strand to help manage CMIP5 files.

GLADE Storage Cluster

- To identify the peculiar delays observed by Dave Mitchell on AIX nodes, examined the packet stats and configuration options on **bluefire**. Identified the disabled SACK (Selective ACK), fasttimo, and NewReno setup on AIX which may not be optimal for a client to Linux glade servers. Helped perform tests on **firefly** and confirmed the improvement in reduction of retransmits under load when sack option is enabled. Recommended that SACK be enabled on the **bluefire** login and I/O nodes.
- The FLEX380 has been flashing the alarm on not being able to use the preferred path. The alarm apparently went off-and-on randomly and there was no clear indication to the failed component. By selectively disabling the services on stratus nodes (GPFS, Lustre, and XFS targets), identified the dependence of alarm triggering events on particular target LUNs and traced the paths over FC for them. The procedure eventually led to the port 18 of fcswitc1 (lower brocade switch) which does not fire up, and migration mirage2_1 (port 18) to gale_p2 (port 1) to resolved the issue. The alarm went away and has stayed off since.
- Upgraded the test gpfs 3.4 on **stratus** nodes from 3.4.0-8 to 3.4.0-13 to prevent possible corruption on mounting older 3.2 based glade file system following after tests on **firefly**.
- Reconfigured the spare LUNs to increase the capacity of /glade/proj3 (had to reboot **oasis4** due to misconfigured cylinder numbers in initial recognition of the new device after zoning change on ddn9900 controllers)
- Had sessions with John Dennis on poor read performance with NetCDF file formats on GPFS. Successfully convinced John Dennis that the poor IO is not due to the file system level performance but the pattern of redundant calls from NCO sub-command, and suggested some temporary options to meet the task deadline.
- Built a new ext3 LUN of 20TB capacity for CMIP5 projects, by reconfiguring the test GPFS3.4 on stratus nodes. Setup a fresh NetCDF package with Cacheless Patch from Hamburg and NCO on mds for tests.
- Enabled the 2TB quotas for all scratch users, minus several exceptions.
- Extended the "Glade Usage Report" Wiki page to show the top 25 users for each filesystem to help CSG monitor scratch users.
- Generated a list of all scratch users with the number and total size of all files written and read on the same day.
- Generated a profile of individual file sizes and totals for all GLADE filesystems.
- Began to document the GPFS policy engine on the DASG wiki.
- Reinstalled GPFS on nomad after filesystem corruption during initial GPFS 3.4 testing.
- Verified the upgrade procedure from GPFS 3.2 to 3.4, especially in regards to the filesystem versions.

